

**NIST Special Publication 1500-9**

---

**NIST Big Data Interoperability  
Framework: Volume 8, Reference  
Architecture Interfaces**

---

NIST Big Data Public Working Group  
Reference Architecture Subgroup

June 2018

This publication is available free of charge from:  
<https://doi.org/10.6028/NIST.SP.1500-9>

**NIST**  
**National Institute of  
Standards and Technology**  
U.S. Department of Commerce

# NIST Special Publication 1500-9

## NIST Big Data Interoperability Framework: Volume 8, Reference Architecture Interfaces

NIST Big Data Public Working Group (NBD-PWG)  
Definitions and Taxonomies Subgroup  
*Information Technology Laboratory*  
*National Institute of Standards and Technology*  
*Gaithersburg, MD 20899*

This publication is available free of charge from:  
<https://doi.org/10.6028/NIST.SP.1500-9>

June 2018



U.S. Department of Commerce  
*Wilbur L. Ross, Jr., Secretary*

National Institute of Standards and Technology  
*Walter Copan, NIST Director and Undersecretary of Commerce for Standards and Technology*

**National Institute of Standards and Technology (NIST) Special Publication 1500-9**  
Natl. Inst. Stand. Technol. Spec. Publ. 1500-9, 92 pages (June 2018) CODEN: NSPUE2

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

NIST Special Publication series 1500 is intended to capture external perspectives related to NIST standards, measurement, and testing-related efforts. These external perspectives can come from industry, academia, government, and others. These reports are intended to document external perspectives and do not represent official NIST positions.

Certain commercial entities, equipment, or materials may be identified in this document to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by NIST, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

There may be references in this publication to other publications currently under development by NIST in accordance with its assigned statutory responsibilities. The information in this publication, including concepts and methodologies, may be used by federal agencies even before the completion of such companion publications. Thus, until each publication is completed, current requirements, guidelines, and procedures, where they exist, remain operative. For planning and transition purposes, federal agencies may wish to closely follow the development of these new publications by NIST.

Organizations are encouraged to review all publications during public comment periods and provide feedback to NIST. All NIST publications are available at <http://www.nist.gov/publication-portal.cfm>.

**Comments on this publication may be submitted to Wo Chang**

National Institute of Standards and Technology  
Attn: Wo Chang, Information Technology Laboratory  
100 Bureau Drive (Mail Stop 8900) Gaithersburg, MD 20899-8930  
Email: [SP1500comments@nist.gov](mailto:SP1500comments@nist.gov)

## Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at NIST promotes the U.S. economy and public welfare by providing technical leadership for the Nation’s measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof of concept implementations, and technical analyses to advance the development and productive use of information technology. ITL’s responsibilities include the development of management, administrative, technical, and physical standards and guidelines for the cost-effective security and privacy of other than national security-related information in federal information systems. This document reports on ITL’s research, guidance, and outreach efforts in Information Technology and its collaborative activities with industry, government, and academic organizations.

### Abstract

This document summarizes interfaces that are instrumental for the interaction with Clouds, Containers, and High Performance Computing (HPC) systems to manage virtual clusters to support the NIST Big Data Reference Architecture (NBDRA). The REpresentational State Transfer (REST) paradigm is used to define these interfaces, allowing easy integration and adoption by a wide variety of frameworks.

Big Data is a term used to describe extensive datasets, primarily in the characteristics of volume, variety, velocity, and/or variability. While opportunities exist with Big Data, the data characteristics can overwhelm traditional technical approaches, and the growth of data is outpacing scientific and technological advances in data analytics. To advance progress in Big Data, the NIST Big Data Public Working Group (NBD-PWG) is working to develop consensus on important fundamental concepts related to Big Data. The results are reported in the *NIST Big Data Interoperability Framework (NBDIF)* series of volumes. This volume, Volume 8, uses the work performed by the NBD-PWG to identify objects instrumental for the NIST Big Data Reference Architecture (NBDRA) which is introduced in the *NBDIF: Volume 6, Reference Architecture*.

### Keywords

Adoption; barriers; implementation; interfaces; market maturity; organizational maturity; project maturity; system modernization.

## Acknowledgements

This document reflects the contributions and discussions by the membership of the NBD-PWG, cochaired by Wo Chang (NIST ITL), Bob Marcus (ET-Strategies), and Chaitan Baru (San Diego Supercomputer Center; National Science Foundation). For all versions, the Subgroups were led by the following people: Nancy Grady (SAIC), Natasha Balac (San Diego Supercomputer Center), and Eugene Luster (R2AD) for the Definitions and Taxonomies Subgroup; Geoffrey Fox (Indiana University) and Tsegereda Beyene (Cisco Systems) for the Use Cases and Requirements Subgroup; Arnab Roy (Fujitsu), Mark Underwood (Krypton Brothers; Synchrony Financial), and Akhil Manchanda (GE) for the Security and Privacy Subgroup; David Boyd (InCadence Strategic Solutions), Orit Levin (Microsoft), Don Krapohl (Augmented Intelligence), and James Ketner (AT&T) for the Reference Architecture Subgroup; and Russell Reinsch (Center for Government Interoperability), David Boyd (InCadence Strategic Solutions), Carl Buffington (Vistrionix), and Dan McClary (Oracle), for the Standards Roadmap Subgroup.

The editors for this document were the following:

- **Version 1:** This volume resulted from Stage 2 work and was not part of the Version 1 scope.
- **Version 2:** Gregor von Laszewski (Indiana University) and Wo Chang (NIST).

Laurie Aldape (Energetics Incorporated) and Elizabeth Lennon (NIST) provided editorial assistance across all NBDIF volumes.

NIST SP 1500-9, NIST Big Data Interoperability Framework: Volume 8, Reference Architecture Interfaces has been collaboratively authored by the NBD-PWG. As of the date of publication, there are over six hundred NBD-PWG participants from industry, academia, and government. Federal agency participants include the National Archives and Records Administration (NARA), National Aeronautics and Space Administration (NASA), National Science Foundation (NSF), and the U.S. Departments of Agriculture, Commerce, Defense, Energy, Census, Health and Human Services, Homeland Security, Transportation, Treasury, and Veterans Affairs.

NIST would like to acknowledge the specific contributions<sup>a</sup> to this volume, during Version 1 and/or Version 2 activities, by the following NBD-PWG members:

**Wo Chang**  
*National Institute of Standard and Technology*

**Geoffrey C. Fox**  
*Indiana University*

**Pratik Thakkar**  
*Philips*

**Gregor von Laszewski**  
*Indiana University*

**Badi Abdhul Wahid**  
*Indiana University*

**Fugang Wang**  
*Indiana University*

**Robert C. Whetsel**  
*DISA/NBIS*

**Alicia Zuniga-Alvarado**  
*Consultant*

<sup>a</sup> “Contributors” are members of the NIST Big Data Public Working Group who dedicated great effort to prepare and gave substantial time on a regular basis to research and development in support of this document.

# TABLE OF CONTENTS

**EXECUTIVE SUMMARY .....IX**

**1 INTRODUCTION ..... 1**

1.1 BACKGROUND..... 1

1.2 SCOPE AND OBJECTIVES OF THE REFERENCE ARCHITECTURES SUBGROUP..... 2

1.3 REPORT PRODUCTION ..... 3

1.4 REPORT STRUCTURE ..... 4

1.5 FUTURE WORK ON THIS VOLUME ..... 4

**2 NBDRA INTERFACE REQUIREMENTS ..... 5**

2.1 HIGH-LEVEL REQUIREMENTS OF THE INTERFACE APPROACH..... 6

2.1.1 *Technology- and Vendor-Agnostic*..... 6

2.1.2 *Support of Plug-In Compute Infrastructure* ..... 6

2.1.3 *Orchestration of Infrastructure and Services*..... 7

2.1.4 *Orchestration of Big Data Applications and Experiments*..... 7

2.1.5 *Reusability* ..... 7

2.1.6 *Execution Workloads* ..... 7

2.1.7 *Security and Privacy Fabric Requirements* ..... 7

2.2 COMPONENT-SPECIFIC INTERFACE REQUIREMENTS..... 8

2.2.1 *System Orchestrator Interface Requirements* ..... 8

2.2.2 *Data Provider Interface Requirements* ..... 8

2.2.3 *Data Consumer Interface Requirements* ..... 9

2.2.4 *Big Data Application Interface Provider Requirements* ..... 9

2.2.4.1 *Collection*..... 9

2.2.4.2 *Preparation*..... 9

2.2.4.3 *Analytics* ..... 10

2.2.4.4 *Visualization* ..... 10

2.2.4.5 *Access* ..... 10

2.2.5 *Big Data Provider Framework Interface Requirements* ..... 10

2.2.5.1 *Infrastructures Interface Requirements* ..... 11

2.2.5.2 *Platforms Interface Requirements* ..... 11

2.2.5.3 *Processing Interface Requirements* ..... 11

2.2.5.4 *Crosscutting Interface Requirements* ..... 11

2.2.5.5 *Messaging/Communications Frameworks* ..... 11

2.2.5.6 *Resource Management Framework*..... 11

**3 SPECIFICATION PARADIGM ..... 13**

3.1 LESSONS LEARNED ..... 13

3.2 HYBRID AND MULTIPLE FRAMEWORKS ..... 13

3.3 DESIGN BY RESEARCH ORIENTED ARCHITECTURE..... 13

3.4 DESIGN BY EXAMPLE..... 13

3.5 INTERFACE COMPLIANCY..... 15

**4 SPECIFICATION..... 16**

4.1 IDENTITY..... 16

4.1.1 *Profile* ..... 16

4.1.2 *User*..... 17

4.1.3 *Organization* ..... 17

4.1.4 *Group/Role*..... 17

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

- 4.2 DATA ..... 18
  - 4.2.1 *TimeStamp* ..... 18
  - 4.2.2 *Variables*..... 19
  - 4.2.3 *Default* ..... 19
  - 4.2.4 *File*..... 20
  - 4.2.5 *Alias* ..... 21
  - 4.2.6 *Replica* ..... 21
  - 4.2.7 *Virtual Directory*..... 22
  - 4.2.8 *Database* ..... 22
  - 4.2.9 *Stream*..... 22
  - 4.2.10 *Filter*..... 22
- 4.3 VIRTUAL CLUSTER..... 23
  - 4.3.1 *Virtual Cluster* ..... 23
  - 4.3.2 *Compute Node*..... 24
  - 4.3.3 *Flavor* ..... 25
  - 4.3.4 *Network Interface Card* ..... 25
  - 4.3.5 *Key*..... 25
  - 4.3.6 *Security Groups* ..... 26
- 4.4 IAAS..... 26
  - 4.4.1 *LibCloud* ..... 26
    - 4.4.1.1 *Challenges*..... 26
    - 4.4.1.2 *LibCloud Flavor*..... 27
    - 4.4.1.3 *LibCloud Image* ..... 27
    - 4.4.1.4 *LibCloud VM*..... 27
    - 4.4.1.5 *LibCloud Node* ..... 28
  - 4.4.2 *OpenStack* ..... 30
    - 4.4.2.1 *OpenStack Flavor* ..... 30
    - 4.4.2.2 *OpenStack Image*..... 30
    - 4.4.2.3 *OpenStack VM* ..... 31
  - 4.4.3 *Azure* ..... 31
    - 4.4.3.1 *Azure Size*..... 32
    - 4.4.3.2 *Azure Image*..... 32
    - 4.4.3.3 *Azure VM* ..... 33
- 4.5 COMPUTE SERVICES..... 33
  - 4.5.1 *Batch Queue*..... 33
  - 4.5.2 *Reservation* ..... 34
- 4.6 CONTAINERS ..... 34
- 4.7 DEPLOYMENT..... 34
- 4.8 MAP/REDUCE..... 35
  - 4.8.1 *Hadoop* ..... 36
- 4.9 MICROSERVICE ..... 37
- 5 STATUS CODES AND ERROR RESPONSES ..... 38**
- APPENDIX A: SCHEMA ..... A-1**
- APPENDIX B: CLOUDMESH REST .....B-1**
  - B.1 PREREQUISITES ..... B-1
  - B.2 REST SERVICE ..... B-1
  - B.3 LIMITATIONS..... B-2
- APPENDIX C: ACRONYMS AND TERMS.....C-1**
- APPENDIX D: BIBLIOGRAPHY ..... D-1**

**FIGURE**

FIGURE 1: NIST BIG DATA REFERENCE ARCHITECTURE (NBDRA) .....5  
 FIGURE 2: NIST BIG DATA REFERENCE ARCHITECTURE INTERFACES .....16  
 FIGURE 3: BOOTING A VM FROM DEFAULTS .....20  
 FIGURE 4: ALLOCATING AND PROVISIONING A VIRTUAL CLUSTER .....23

**TABLE**

TABLE 1: HTTP RESPONSE CODES .....38

**OBJECTS**

OBJECT 3.1: EXAMPLE OBJECT SPECIFICATION .....14  
 OBJECT 4.1: PROFILE .....16  
 OBJECT 4.2: USER .....17  
 OBJECT 4.3: ORGANIZATION .....17  
 OBJECT 4.4: GROUP .....17  
 OBJECT 4.5: ROLE .....18  
 OBJECT 4.6: TIMESTAMP .....19  
 OBJECT 4.7: VARIABLES .....19  
 OBJECT 4.8: DEFAULT .....19  
 OBJECT 4.9: FILE .....21  
 OBJECT 4.10: FILE ALIAS .....21  
 OBJECT 4.11: REPLICATION .....21  
 OBJECT 4.12: VIRTUAL DIRECTORY .....22  
 OBJECT 4.13: DATABASE .....22  
 OBJECT 4.14: STREAM .....22  
 OBJECT 4.15: FILTER .....23  
 OBJECT 4.16: VIRTUAL CLUSTER .....24  
 OBJECT 4.17: VIRTUAL CLUSTER PROVIDER .....24  
 OBJECT 4.18: COMPUTE NODE OF A VIRTUAL CLUSTER .....24  
 OBJECT 4.19: FLAVOR .....25  
 OBJECT 4.20: NETWORK INTERFACE CARD .....25  
 OBJECT 4.21: KEY .....25  
 OBJECT 4.22: SECURITY GROUPS .....26  
 OBJECT 4.23: LIBCLOUD FLAVOR .....27  
 OBJECT 4.24: LIBCLOUD IMAGE .....27  
 OBJECT 4.25: LIBCLOUD VM .....28  
 OBJECT 4.26: LIBCLOUD NODE .....28  
 OBJECT 4.27: OPENSTACK FLAVOR .....30  
 OBJECT 4.28: OPENSTACK IMAGE .....30  
 OBJECT 4.29: OPENSTACK VM .....31  
 OBJECT 4.30: AZURE-SIZE .....32  
 OBJECT 4.31: AZURE-IMAGE .....32  
 OBJECT 4.32: AZURE VM .....33  
 OBJECT 4.33: BATCH JOB .....33  
 OBJECT 4.34: RESERVATION .....34  
 OBJECT 4.35: CONTAINER .....34  
 OBJECT 4.36: DEPLOYMENT .....34  
 OBJECT 4.37: MAP/REDUCE .....35

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>



OBJECT 4.38: MAP/REDUCE FUNCTION ..... 36  
OBJECT 4.39: MAP/REDUCE NOOP..... 36  
OBJECT 4.40: HADOOP ..... 37  
OBJECT 4.41: MICROSERVICE..... 37  
OBJECT A.1: SCHEMA..... A-1

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

# EXECUTIVE SUMMARY

---

The *NIST Big Data Interoperability Framework (NBDIF): Volume 8, Reference Architecture Interfaces* document was prepared by the NIST Big Data Public Working Group (NBD-PWG) Interface Subgroup to identify interfaces in support of the NIST Big Data Reference Architecture (NBDRA). The interfaces contain two different aspects:

- The definition of resources that are part of the NBDRA. These resources are formulated in JSON format and can be integrated into a REST framework or an object-based framework easily.
- The definition of simple interface use cases that allow us to illustrate the usefulness of the resources defined.

The resources were categorized in groups that are identified by the NBDRA set forward in the *NBDIF: Volume 6, Reference Architecture* document. While the *NBDIF: Volume 3, Use Cases and General Requirements* document provides *application-oriented* high-level use cases, the use cases defined in this document are subsets of them and focus on *interface* use cases. The interface use cases are not meant to be complete examples, but showcase why the resource has been defined. Hence, the interfaces use cases are only representative, and do not represent the entire spectrum of Big Data usage. All the interfaces were openly discussed in the working group. Additions are welcome and we like to discuss your contributions in the group.

The NIST Big Data Interoperability Framework consists of nine volumes, each of which addresses a specific key topic, resulting from the work of the NBD-PWG. The nine volumes are:

- Volume 1: Definitions [1]
- Volume 2: Taxonomies [2]
- Volume 3: Use Cases and General Requirements [3]
- Volume 4: Security and Privacy [4]
- Volume 5: Architectures White Paper Survey [5]
- Volume 6: Reference Architecture [6]
- Volume 7: Standards Roadmap [7]
- Volume 8: Reference Architecture Interfaces (this volume)
- Volume 9: Adoption and Modernization [8]

The NBDIF will be released in three versions, which correspond to the three development stages of the NBD-PWG work. The three stages aim to achieve the following with respect to the NBDRA.

Stage 1: Identify the high-level Big Data reference architecture key components, which are technology-, infrastructure-, and vendor-agnostic.

Stage 2: Define general interfaces between the NBDRA components.

Stage 3: Validate the NBDRA by building Big Data general applications through the general interfaces.

This document is targeting Stage 2 of the NBDRA. Coordination of the group is conducted on the NBD-PWG web page (<https://bigdatawg.nist.gov>).

# 1 INTRODUCTION

---

## 1.1 BACKGROUND

There is broad agreement among commercial, academic, and government leaders about the remarkable potential of Big Data to spark innovation, fuel commerce, and drive progress. Big Data is the common term used to describe the deluge of data in today's networked, digitized, sensor-laden, and information-driven world. The availability of vast data resources carries the potential to answer questions previously out of reach, including the following:

- How can a potential pandemic reliably be detected early enough to intervene?
- Can new materials with advanced properties be predicted before these materials have ever been synthesized?
- How can the current advantage of the attacker over the defender in guarding against cybersecurity threats be reversed?

There is also broad agreement on the ability of Big Data to overwhelm traditional approaches. The growth rates for data volumes, speeds, and complexity are outpacing scientific and technological advances in data analytics, management, transport, and data user spheres.

Despite widespread agreement on the inherent opportunities and current limitations of Big Data, a lack of consensus on some important fundamental questions continues to confuse potential users and stymie progress. These questions include the following:

- How is Big Data defined?
- What attributes define Big Data solutions?
- What is new in Big Data?
- What is the difference between Big Data and *bigger data* that has been collected for years?
- How is Big Data different from traditional data environments and related applications?
- What are the essential characteristics of Big Data environments?
- How do these environments integrate with currently deployed architectures?
- What are the central scientific, technological, and standardization challenges that need to be addressed to accelerate the deployment of robust, secure Big Data solutions?

Within this context, on March 29, 2012, the White House announced the Big Data Research and Development Initiative.<sup>1</sup> The initiative's goals include helping to accelerate the pace of discovery in science and engineering, strengthening national security, and transforming teaching and learning by improving analysts' ability to extract knowledge and insights from large and complex collections of digital data.

Six federal departments and their agencies announced more than \$200 million in commitments spread across more than 80 projects, which aim to significantly improve the tools and techniques needed to access, organize, and draw conclusions from huge volumes of digital data. The initiative also challenged industry, research universities, and nonprofits to join with the federal government to make the most of the opportunities created by Big Data.

Motivated by the White House initiative and public suggestions, the National Institute of Standards and Technology (NIST) accepted the challenge to stimulate collaboration among industry professionals to further the secure and effective adoption of Big Data. As a result of NIST's Cloud and Big Data Forum held on January 15–17, 2013, there was strong encouragement for NIST to create a public working group for the development of a Big Data Standards Roadmap. Forum participants noted that this roadmap should define and prioritize Big Data

requirements, including interoperability, portability, reusability, extensibility, data usage, analytics, and technology infrastructure. In doing so, the roadmap would accelerate the adoption of the most secure and effective Big Data techniques and technology.

On June 19, 2013, the NIST Big Data Public Working Group (NBD-PWG) was launched with extensive participation by industry, academia, and government from across the nation. The scope of the NBD-PWG involves forming a community of interests from all sectors—including industry, academia, and government—with the goal of developing consensus on definitions, taxonomies, secure reference architectures, security and privacy, and, from these, a standards roadmap. Such a consensus would create a vendor-neutral, technology- and infrastructure-independent framework that would enable Big Data stakeholders to identify and use the best analytics tools for their processing and visualization requirements on the most suitable computing platform and cluster, while also allowing added value from Big Data service providers.

The *NIST Big Data Interoperability Framework* (NBDIF) will be released in three versions, which correspond to the three stages of the NBD-PWG work. The three stages aim to achieve the following with respect to the NIST Big Data Reference Architecture (NBDRA).

Stage 1: Identify the high-level Big Data reference architecture key components, which are technology, infrastructure, and vendor agnostic.

Stage 2: Define general interfaces between the NBDRA components.

Stage 3: Validate the NBDRA by building Big Data general applications through the general interfaces.

On September 16, 2015, seven NBDIF Version 1 volumes were published ([http://bigdatawg.nist.gov/V1\\_output\\_docs.php](http://bigdatawg.nist.gov/V1_output_docs.php)), each of which addresses a specific key topic, resulting from the work of the NBD-PWG. The seven volumes are as follows:

- Volume 1, Definitions [1]
- Volume 2, Taxonomies [2]
- Volume 3, Use Cases and General Requirements [3]
- Volume 4, Security and Privacy [4]
- Volume 5, Architectures White Paper Survey [5]
- Volume 6, Reference Architecture [6]
- Volume 7, Standards Roadmap [7]

Currently, the NBD-PWG is working on Stage 2 with the goals to enhance the Version 1 content, define general interfaces between the NBDRA components by aggregating low-level interactions into high-level general interfaces, and demonstrate how the NBDRA can be used. As a result of the Stage 2 work, the following two additional NBDIF volumes have been developed.

- Volume 8, Reference Architecture Interfaces (this volume)
- Volume 9, Adoption and Modernization [8]

Version 2 of the NBDIF volumes, resulting from Stage 2 work, can be downloaded from the NBD-PWG website ([https://bigdatawg.nist.gov/V2\\_output\\_docs.php](https://bigdatawg.nist.gov/V2_output_docs.php)). Potential areas of future work for each volume during Stage 3 are highlighted in Section 1.5 of each volume. The current effort documented in this volume reflects concepts developed within the rapidly evolving field of Big Data.

## 1.2 SCOPE AND OBJECTIVES OF THE REFERENCE ARCHITECTURES SUBGROUP

Reference architectures provide “an authoritative source of information about a specific subject area that guides and constrains the instantiations of multiple architectures and solutions.” Reference architectures generally serve as a foundation for solution architectures and may also be used for comparison and alignment of instantiations of architectures and solutions.

The goal of the NBD-PWG Reference Architecture Subgroup is to develop an open reference architecture for Big Data that achieves the following objectives:

- Provides a common language for the various stakeholders;
- Encourages adherence to common standards, specifications, and patterns;
- Provides consistent methods for implementation of technology to solve similar problem sets;
- Illustrates and improves understanding of the various Big Data components, processes, and systems, in the context of a vendor- and technology-agnostic Big Data conceptual model;
- Provides a technical reference for U.S. government departments, agencies, and other consumers to understand, discuss, categorize, and compare Big Data solutions; and
- Facilitates analysis of candidate standards for interoperability, portability, reusability, and extensibility.

The NBDRA is a high-level conceptual model crafted to serve as a tool to facilitate open discussion of the requirements, design structures, and operations inherent in Big Data. The NBDRA is intended to facilitate the understanding of the operational intricacies in Big Data. It does not represent the system architecture of a specific Big Data system, but rather is a tool for describing, discussing, and developing system-specific architectures using a common framework of reference. The model is not tied to any specific vendor products, services, or reference implementation, nor does it define prescriptive solutions that inhibit innovation.

The NBDRA does not address the following:

- Detailed specifications for any organization's operational systems;
- Detailed specifications of information exchanges or services; and
- Recommendations or standards for integration of infrastructure products.

The goals of the Subgroup will be realized throughout the three planned phases of the NBD-PWG work, as outlined in Section 1.1.

### 1.3 REPORT PRODUCTION

The *NBDIF: Volume 8, References Architecture Interfaces* is one of nine volumes, whose overall aims are to define and prioritize Big Data requirements, including interoperability, portability, reusability, extensibility, data usage, analytic techniques, and technology infrastructure to support secure and effective adoption of Big Data. The overall goals of this volume are to define and specify interfaces to implement the Big Data Reference Architecture. This volume arose from discussions during the weekly NBD-PWG conference calls. Topics included in this volume began to take form in Phase 2 of the NBD-PWG work. This volume represents the groundwork for additional content planned for Phase 3. During the discussions, the NBD-PWG identified the need to specify a variety of interfaces.

To enable interoperability between the NBDRA components, a list of well-defined NBDRA interfaces is needed. These interfaces are documented in this volume. To introduce them, the NBDRA structure will be followed, focusing on interfaces that allow bootstrapping of the NBDRA. The document begins with a summary of requirements that will be integrated into our specifications. Subsequently, each section will introduce a number of objects that build the core of the interface addressing a specific aspect of the NBDRA. A selected number of *interface use cases* will be showcased to outline how the specific interface can be used in a reference implementation of the NBDRA. Validation of this approach can be achieved while applying it to the application use cases that have been gathered in the *NBDIF: Volume 3, Use Cases and Requirements* document.

These application use cases have considerably contributed towards the design of the NBDRA. Hence the expectation for the interfaces is to enable implementation of NBDRA for the existing or new use cases. This approach can facilitate subsequent analysis and comparison of the use cases.

This document is expected to grow with the help of contributions from the community to achieve a comprehensive set of interfaces that will enable developers to build any Big Data applications based from the NBDRA Big Data Architecture. To achieve technical and high-quality document content, this document will go through a public comment period along with NIST internal review.

## 1.4 REPORT STRUCTURE

The organization of this document roughly corresponds to the process used by the NBD-PWG to develop the interfaces. Following the introductory material presented in Section 1, the remainder of this document is organized as follows:

- Section 2 presents the interface requirements;
- Section 3 summarizes the elementary objects that are important to the NBDRA;
- Section 4 presents several objects grouped by functional use; and
- Four appendices provide supplementary information.

## 1.5 FUTURE WORK ON THIS VOLUME

A number of topics have not been discussed and clarified sufficiently to be included in Version 2. Version 3 topics will be identified during discussions within the Reference Architecture Subgroup.

## 2 NBDRA INTERFACE REQUIREMENTS

The development of a Big Data reference architecture requires a thorough understanding of current techniques, issues, and concerns. To this end, the NBD-PWG collected use cases to gain an understanding of current applications of Big Data, conducted a survey of reference architectures to understand commonalities within Big Data architectures in use, developed a taxonomy to understand and organize the information collected, and reviewed existing technologies and trends relevant to Big Data. The results of these NBD-PWG activities were used in the development of the NBDRA (Figure 1) and the interfaces presented herein. Detailed descriptions of these activities can be found in the other volumes of the *NBDIF*.

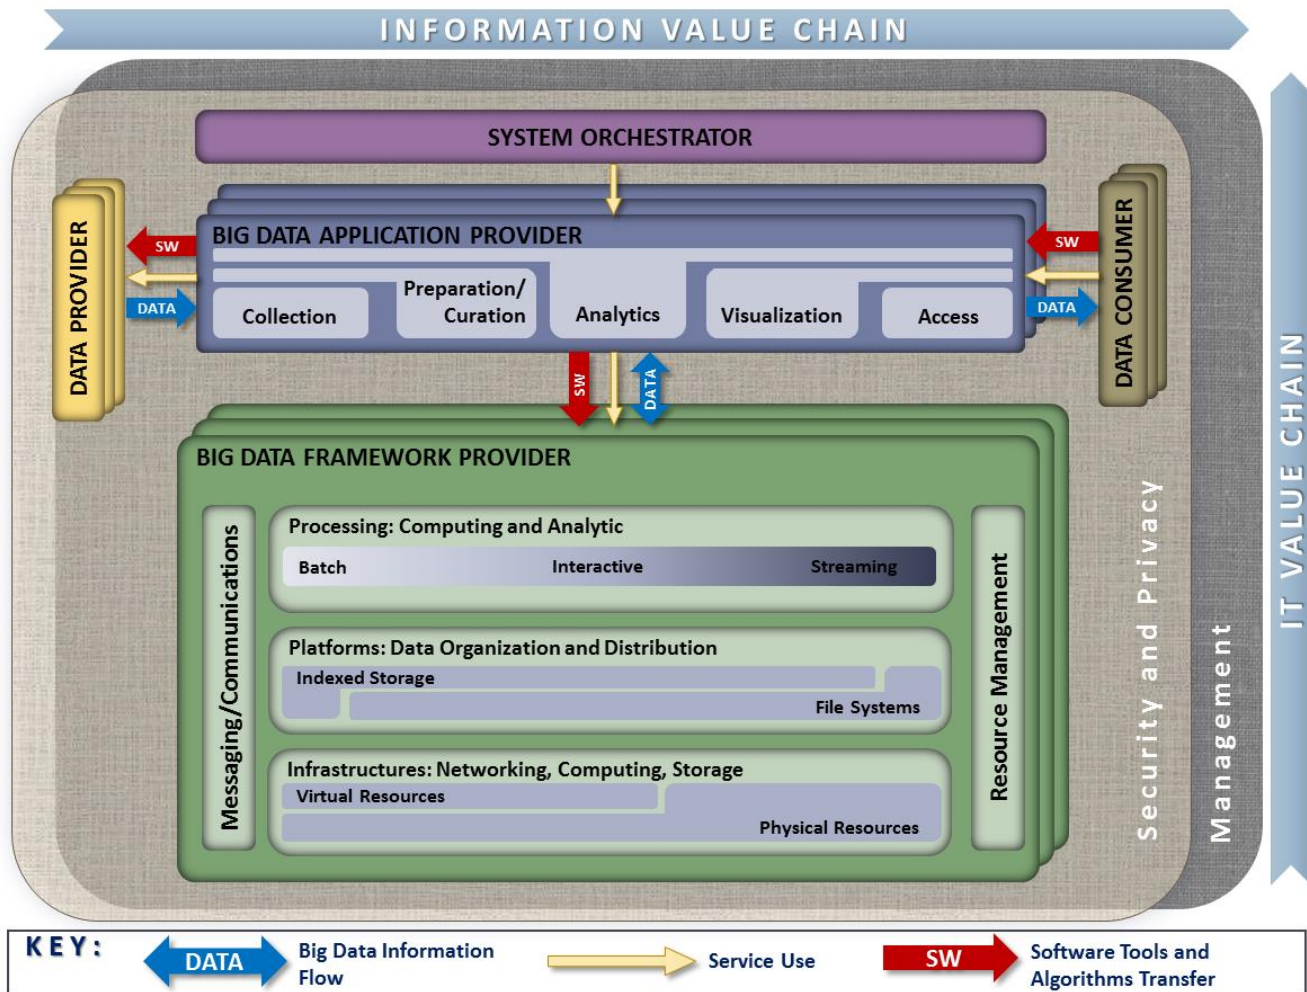


Figure 1: NIST Big Data Reference Architecture (NBDRA)

This vendor-neutral, technology- and infrastructure-agnostic conceptual model, the NBDRA, is shown in Figure 1 and represents a Big Data system composed of five logical functional components connected by interoperability interfaces (i.e., services). Two fabrics envelop the components, representing the interwoven nature of management and security and privacy with all five of the components. These two fabrics provide services and functionality to the five main roles in the areas specific to Big Data and are crucial to any Big Data solution. Note: None of the terminology or diagrams in these documents is intended to be normative or to imply any business or

deployment model. The terms *provider* and *consumer* as used are descriptive of general roles and are meant to be informative in nature.

The NBDRA is organized around five major roles and multiple sub-roles aligned along two axes representing the two Big Data value chains: Information Value (horizontal axis) and Information Technology (IT; vertical axis). Along the Information Value axis, the value is created by data collection, integration, analysis, and applying the results following the value chain. Along the IT axis, the value is created by providing networking, infrastructure, platforms, application tools, and other IT services for hosting of and operating the Big Data in support of required data applications. At the intersection of both axes is the Big Data Application Provider role, indicating that data analytics and its implementation provide the value to Big Data stakeholders in both value chains. The term *provider* as part of the Big Data Application Provider and Big Data Framework Provider is there to indicate that those roles provide or implement specific activities and functions within the system. It does not designate a service model or business entity.

The DATA arrows in Figure 2 show the flow of data between the system's main roles. Data flows between the roles either physically (i.e., by value) or by providing its location and the means to access it (i.e., by reference). The SW arrows show transfer of software tools for processing of Big Data *in situ*. The Service Use arrows represent software programmable interfaces. While the main focus of the NBDRA is to represent the run-time environment, all three types of communications or transactions can happen in the configuration phase as well. Manual agreements (e.g., service-level agreements) and human interactions that may exist throughout the system are not shown in the NBDRA.

Detailed information on the NBDRA conceptual model is presented in the *NBDIF: Volume 6, Reference Architecture* document. [6]

Prior to outlining the specific interfaces, general requirements are introduced and the interfaces are defined.

## 2.1 HIGH-LEVEL REQUIREMENTS OF THE INTERFACE APPROACH

This section focuses on the high-level requirements of the interface approach that are needed to implement the reference architecture depicted in Figure 1.

### 2.1.1 TECHNOLOGY- AND VENDOR-AGNOSTIC

Due to the many different tools, services, and infrastructures available in the general area of Big Data, an interface ought to be as vendor-independent as possible, while at the same time be able to leverage best practices. Hence, a methodology is needed that allows extension of interfaces to adapt and leverage existing approaches, but also allows the interfaces to provide merit in easy specifications that assist the formulation and definition of the NBDRA.

### 2.1.2 SUPPORT OF PLUG-IN COMPUTE INFRASTRUCTURE

As Big Data is not just about hosting data, but about analyzing data, the interfaces provided herein must encapsulate a rich infrastructure environment that is used by data scientists. This includes the ability to integrate (or plug-in) various compute resources and services to provide the necessary compute power to analyze the data. These resources and services include the following:

- Access to hierarchy of compute resources from the laptop/desktop, servers, data clusters, and clouds;
- The ability to integrate special-purpose hardware such as GPUs and FPGAs that are used in accelerated analysis of data; and
- The integration of services including micro services that allow the analysis of the data by delegating them to hosted or dynamically deployed services on the infrastructure of choice.



### 2.1.3 ORCHESTRATION OF INFRASTRUCTURE AND SERVICES

From review of the use case collection, presented in the *NBDIF: Volume 3, Use Cases and General Requirements* document [3], the need arose to address the mechanism of preparing suitable infrastructures for various use cases. As not every infrastructure is suited for every use case, a custom infrastructure may be needed. As such, this document is not attempting to deliver a single deployed NBDRA, but allow the setup of an infrastructure that satisfies the particular use case. To achieve this task, it is necessary to provision software stacks and services while orchestrating their deployment and leveraging infrastructures. It is not the focus of this document to replace existing orchestration software and services, but provide an interface to them to leverage them as part of defining and creating the infrastructure. Various orchestration frameworks and services could therefore be leveraged, even as part of the same framework, and work in orchestrated fashion to achieve the goal of preparing an infrastructure suitable for one or more applications.

### 2.1.4 ORCHESTRATION OF BIG DATA APPLICATIONS AND EXPERIMENTS

The creation of the infrastructure suitable for Big Data applications provides the basic computing environment. However, Big Data applications may require the creation of sophisticated applications as part of interactive experiments to analyze and probe the data. For this purpose, the applications must be able to orchestrate and interact with experiments conducted on the data while assuring reproducibility and correctness of the data. For this purpose, a *System Orchestrator* (either the data scientists or a service acting on behalf of the data scientist) is used as the command center to interact on behalf of the Big Data Application Provider to orchestrate dataflow from Data Provider, carry out the Big Data application life cycle with the help of the Big Data Framework Provider, and enable the Data Consumer to consume Big Data processing results. An interface is needed to describe these interactions and to allow leveraging of experiment management frameworks in scripted fashion. A customization of parameters is needed on several levels. On the highest level, high-level, application-motivated parameters are needed to drive the orchestration of the experiment. On lower levels, these high-level parameters may drive and create service-level agreements, augmented specifications, and parameters that could even lead to the orchestration of infrastructure and services to satisfy experiment needs.

### 2.1.5 REUSABILITY

The interfaces provided must encourage reusability of the infrastructure, services, and experiments described by them. This includes (1) reusability of available analytics packages and services for adoption; (2) deployment of customizable analytics tools and services; and (3) operational adjustments that allow the services and infrastructure to be adapted while at the same time allowing for reproducible experiment execution.

### 2.1.6 EXECUTION WORKLOADS

One of the important aspects of distributed Big Data services can be that the data served is simply too big to be moved to a different location. Instead, an interface could allow the description and packaging of analytics algorithms, and potentially also tools, as a payload to a data service. This can be best achieved, not by sending the detailed execution, but by sending an interface description that describes how such an algorithm or tool can be created on the server and be executed under security considerations (integrated with authentication and authorization in mind).

### 2.1.7 SECURITY AND PRIVACY FABRIC REQUIREMENTS

Although the focus of this document is not security and privacy, which are documented in the *NBDIF: Volume 4, Security and Privacy* [4], the interfaces defined herein must be capable of integration into a secure reference architecture that supports secure execution, secure data transfer, and privacy. Consequently, the interfaces defined herein can be augmented with frameworks and solutions that provide such mechanisms. Thus, diverse requirement needs stemming from different use cases addressing security need to be distinguished. To contrast that the security requirements between applications can vary drastically, the following example is provided.

Although many of the interfaces and their objects to support Big Data applications in physics are similar to those in healthcare, they differ in the integration of security interfaces and policies. While in physics the protection of data is less of an issue, it is a stringent requirement in healthcare. Thus, deriving architectural frameworks for both may use largely similar components, but addressing security issues will be very different. In future versions of this document, the security of interfaces may be addressed. In the meanwhile, they are considered an advanced use case showcasing that the validity of the specifications introduced here is preserved, even if security and privacy requirements differ vastly among application use cases.

## 2.2 COMPONENT-SPECIFIC INTERFACE REQUIREMENTS

This section summarizes the requirements for the interfaces of the NBDRA components. The five components are listed in Figure 1 and addressed in each of the subsections as part of Section 2.2.1 (System Orchestrator Interface Requirements) and Section 2.2.6 (Big Data Application Provider to Big Data Framework Provider Interface) of this document. The five main functional components of the NBDRA represent the different technical roles within a Big Data system. The functional components are listed below and discussed in subsequent subsections.

- System Orchestrator: Defines and integrates the required data application activities into an operational vertical system (see Section 2.2.1);
- Data Provider: Introduces new data or information feeds into the Big Data system (see Section 2.2.2);
- Data Consumer: Includes end users or other systems that use the results of the Big Data Application Provider (see Section 2.2.3);
- Big Data Application Provider: Executes a data life cycle to meet security and privacy requirements as well as System Orchestrator-defined requirements (see Section 2.2.4);
- Big Data Framework Provider: Establishes a computing framework in which to execute certain transformation applications while protecting the privacy and integrity of data (see Section 2.2.5); and

### 2.2.1 SYSTEM ORCHESTRATOR INTERFACE REQUIREMENTS

The System Orchestrator role includes defining and integrating the required data application activities into an operational vertical system. Typically, the System Orchestrator involves a collection of more specific roles, performed by one or more actors, which manage and orchestrate the operation of the Big Data system. These actors may be human components, software components, or some combination of the two. The function of the System Orchestrator is to configure and manage the other components of the Big Data architecture to implement one or more workloads that the architecture is designed to execute. The workloads managed by the System Orchestrator may be assigning/provisioning framework components to individual physical or virtual nodes at the lower level, or providing a graphical user interface that supports the specification of workflows linking together multiple applications and components at the higher level. The System Orchestrator may also, through the Management Fabric, monitor the workloads and system to confirm that specific quality of service requirements is met for each workload, and may elastically assign and provision additional physical or virtual resources to meet workload requirements resulting from changes/surges in the data or number of users/transactions. The interface to the System Orchestrator must be capable of specifying the task of orchestration the deployment, configuration, and the execution of applications within the NBDRA. A simple vendor-neutral specification to coordinate the various parts either as simple parallel language tasks or as a workflow specification is needed to facilitate the overall coordination. Integration of existing tools and services into the System Orchestrator as extensible interfaces is desirable.

### 2.2.2 DATA PROVIDER INTERFACE REQUIREMENTS

The Data Provider role introduces new data or information feeds into the Big Data system for discovery, access, and transformation by the Big Data system. New data feeds are distinct from the data already in use by the system and residing in the various system repositories. Similar technologies can be used to access both new data feeds and existing data. The Data Provider actors can be anything from a sensor, to a human inputting data manually, to

another Big Data system. Interfaces for data providers must be able to specify a data provider so it can be located by a data consumer. It also must include enough details to identify the services offered so they can be pragmatically reused by consumers. Interfaces to describe pipes and filters must be addressed.

### **2.2.3 DATA CONSUMER INTERFACE REQUIREMENTS**

Like the Data Provider, the role of Data Consumer within the NBDRA can be an actual end user or another system. In many ways, this role is the mirror image of the Data Provider, with the entire Big Data framework appearing like a Data Provider to the Data Consumer. The activities associated with the Data Consumer role include the following:

- Search and Retrieve,
- Download,
- Analyze Locally,
- Reporting,
- Visualization, and
- Data to Use for Their Own Processes.

The interface for the data consumer must be able to describe the consuming services and how they retrieve information or leverage data consumers.

### **2.2.4 BIG DATA APPLICATION INTERFACE PROVIDER REQUIREMENTS**

The Big Data Application Provider role executes a specific set of operations along the data life cycle to meet the requirements established by the System Orchestrator, as well as meeting security and privacy requirements. The Big Data Application Provider is the architecture component that encapsulates the business logic and functionality to be executed by the architecture. The interfaces to describe Big Data applications include interfaces for the various subcomponents including collections, preparation/curation, analytics, visualization, and access. Some of the interfaces used in these subcomponents can be reused from other interfaces, which are introduced in other sections of this document. Where appropriate, application-specific interfaces will be identified and examples provided with a focus on use cases as identified in the *NBDIF: Volume 3, Use Cases and General Requirements*.

#### **2.2.4.1 Collection**

In general, the collection activity of the Big Data Application Provider handles the interface with the Data Provider. This may be a general service, such as a file server or web server configured by the System Orchestrator to accept or perform specific collections of data, or it may be an application-specific service designed to pull data or receive pushes of data from the Data Provider. Since this activity is receiving data at a minimum, it must store/buffer the received data until it is persisted through the Big Data Framework Provider. This persistence need not be to physical media but may simply be to an in-memory queue or other service provided by the processing frameworks of the Big Data Framework Provider. The collection activity is likely where the extraction portion of the Extract, Transform, Load (ETL)/Extract, Load, Transform (ELT) cycle is performed. At the initial collection stage, sets of data (e.g., data records) of similar structure are collected (and combined), resulting in uniform security, policy, and other considerations. Initial metadata is created (e.g., subjects with keys are identified) to facilitate subsequent aggregation or look-up methods.

#### **2.2.4.2 Preparation**

The preparation activity is where the transformation portion of the ETL/ELT cycle is likely performed, although analytics activity will also likely perform advanced parts of the transformation. Tasks performed by this activity could include data validation (e.g., checksums/hashes, format checks), cleansing (e.g., eliminating bad records/fields), outlier removal, standardization, reformatting, or encapsulating. This activity is also where source data will frequently be persisted to archive storage in the Big Data Framework Provider and provenance data will be verified or attached/associated. Verification or attachment may include optimization of data through

manipulations (e.g., deduplication) and indexing to optimize the analytics process. This activity may also aggregate data from different Data Providers, leveraging metadata keys to create an expanded and enhanced dataset.

#### **2.2.4.3 Analytics**

The analytics activity of the Big Data Application Provider includes the encoding of the low-level business logic of the Big Data system (with higher-level business process logic being encoded by the System Orchestrator). The activity implements the techniques to extract knowledge from the data based on the requirements of the vertical application. The requirements specify the data processing algorithms to produce new insights that will address the technical goal. The analytics activity will leverage the processing frameworks to implement the associated logic. This typically involves the activity providing software that implements the analytic logic to the batch and/or streaming elements of the processing framework for execution. The messaging/communication framework of the Big Data Framework Provider may be used to pass data or control functions to the application logic running in the processing frameworks. The analytic logic may be broken up into multiple modules to be executed by the processing frameworks which communicate, through the messaging/communication framework, with each other and other functions instantiated by the Big Data Application Provider.

#### **2.2.4.4 Visualization**

The visualization activity of the Big Data Application Provider prepares elements of the processed data and the output of the analytic activity for presentation to the Data Consumer. The objective of this activity is to format and present data in such a way as to optimally communicate meaning and knowledge. The visualization preparation may involve producing a text-based report or rendering the analytic results as some form of graphic. The resulting output may be a static visualization and may simply be stored through the Big Data Framework Provider for later access. However, the visualization activity frequently interacts with the access activity, the analytics activity, and the Big Data Framework Provider (processing and platform) to provide interactive visualization of the data to the Data Consumer based on parameters provided to the access activity by the Data Consumer. The visualization activity may be completely application-implemented, leverage one or more application libraries, or may use specialized visualization processing frameworks within the Big Data Framework Provider.

#### **2.2.4.5 Access**

The access activity within the Big Data Application Provider is focused on the communication/interaction with the Data Consumer. Like the collection activity, the access activity may be a generic service such as a web server or application server that is configured by the System Orchestrator to handle specific requests from the Data Consumer. This activity would interface with the visualization and analytic activities to respond to requests from the Data Consumer (who may be a person) and uses the processing and platform frameworks to retrieve data to respond to Data Consumer requests. In addition, the access activity confirms that descriptive and administrative metadata and metadata schemes are captured and maintained for access by the Data Consumer and as data is transferred to the Data Consumer. The interface with the Data Consumer may be synchronous or asynchronous in nature and may use a pull or push paradigm for data transfer.

### **2.2.5 BIG DATA PROVIDER FRAMEWORK INTERFACE REQUIREMENTS**

Data for Big Data applications are delivered through data providers. They can be either local providers, data contributed by a user, or distributed data providers, data on the Internet. This interface must be able to provide the following functionality:

- Interfaces to files,
- Interfaces to virtual data directories,
- Interfaces to data streams, and
- Interfaces to data filters.

The Big Data Framework Provider typically consists of one or more hierarchically organized instances of the components in the NBDRA IT value chain (Figure 1). There is no requirement that all instances at a given level in the hierarchy be of the same technology. In fact, most Big Data implementations are hybrids that combine multiple technology approaches to provide flexibility or meet the complete range of requirements, which are driven from the Big Data Application Provider.

### **2.2.5.1 Infrastructures Interface Requirements**

This Big Data Framework Provider element provides all the resources necessary to host/run the activities of the other components of the Big Data system. Typically, these resources consist of some combination of physical resources, which may host/support similar virtual resources. The NBDRA needs interfaces that can be used to deal with the underlying infrastructure to address networking, computing, and storage.

### **2.2.5.2 Platforms Interface Requirements**

As part of the NBDRA platforms, interfaces are needed that can address platform needs and services for data organization, data distribution, indexed storage, and file systems.

### **2.2.5.3 Processing Interface Requirements**

The processing frameworks for Big Data provide the necessary infrastructure software to support implementation of applications that can deal with the volume, velocity, variety, and variability of data. Processing frameworks define how the computation and processing of the data is organized. Big Data applications rely on various platforms and technologies to meet the challenges of scalable data analytics and operation. A requirement is the ability to interface easily with computing services that offer specific analytics services, batch processing capabilities, interactive analysis, and data streaming.

### **2.2.5.4 Crosscutting Interface Requirements**

Several crosscutting interface requirements within the Big Data Framework Provider include messaging, communication, and resource management. Often these services may be hidden from explicit interface use as they are part of larger systems that expose higher-level functionality through their interfaces. However, such interfaces may also be exposed on a lower level in case finer-grained control is needed. The need for such crosscutting interface requirements will be extracted from the *NBDIF: Volume 3, Use Cases and General Requirements* document.

### **2.2.5.5 Messaging/Communications Frameworks**

Messaging and communications frameworks have their roots in the High Performance Computing environments long popular in the scientific and research communities. Messaging/Communications Frameworks were developed to provide application programming interfaces (APIs) for the reliable queuing, transmission, and receipt of data.

### **2.2.5.6 Resource Management Framework**

As Big Data systems have evolved and become more complex, and as businesses work to leverage limited computation and storage resources to address a broader range of applications and business challenges, the requirement to effectively manage those resources has grown significantly. While tools for resource management and *elastic computing* have expanded and matured in response to the needs of cloud providers and virtualization

technologies, Big Data introduces unique requirements for these tools. However, Big Data frameworks tend to fall more into a distributed computing paradigm, which presents additional challenges.

## 3 SPECIFICATION PARADIGM

---

This section summarizes the elementary objects that are important to the NBDRA.

### 3.1 LESSONS LEARNED

Originally, a full REpresentational State Transfer (REST) specification was used for defining the objects related to the NBDRA [9]. However, at this stage of the document, it would introduce too complex a notation framework. This would result in (1) a considerable increase in length of this document; (2) a more complex framework reducing potential public participation in the project; and (3) a more complex framework for developing a reference implementation. Thus, in this version of the document, a design concept by example will be introduced, which is used to automatically create a schema as well as a reference implementation.

### 3.2 HYBRID AND MULTIPLE FRAMEWORKS

To avoid vendor lock-in, Big Data systems must be able to deal with hybrid and multiple frameworks. This is not only true for Clouds, containers, DevOps, but also for components of the NBDRA.

### 3.3 DESIGN BY RESEARCH ORIENTED ARCHITECTURE

A resource-oriented architecture represents a software architecture and programming paradigm for designing and developing software in the form of resources. It is often associated with *RESTful* interfaces. The resources are software components which can be reused in concrete reference implementations.

### 3.4 DESIGN BY EXAMPLE

To accelerate discussion among the NBD-PWG members, an approach by example is used to define objects and their interfaces. These examples can then be used to automatically generate a schema. The schema is added to Appendix A of the document. Appendix A lists the schema that is automatically created from the definitions. More information about the creation can be found in Appendix B.

Focusing first on examples allows acceleration of the design process and simplification of discussions about the objects and interfaces. Hence, complex specifications should not encumber the development. The processes and specifications used in this document will also permit automatic creation of an implementation of the objects that can be integrated into a reference architecture, such as the cloudmesh client and REST project [10] [9].

An example object will demonstrate this approach. The following object defines a JSON object representing a user.

**Object 3.1: Example Object Specification**

```
{
  "profile": {
    "description": "The Profile of a user",
    "uuid": "jshdjkdh...",
    "context": "resource",
    "email": "laszewski@gmail.com",
    "firstname": "Gregor",
    "lastname": "von Laszewski",
    "username": "gregor",
    "publickey": "ssh ...."
  }
}
```

Such an object can be translated to a schema specification while introspecting the types of the original example.

All examples are managed in Github and links to them are automatically generated to be included into this document. The resulting schema object follows the Cerberus [11] specification and looks for the specific object, introduced earlier, as follows:

```
profile = {
  'schema': {
    'username': {'type': 'string'},
    'context': {'type': 'string'},
    'description': {'type': 'string'},
    'firstname': {'type': 'string'},
    'lastname': {'type': 'string'},
    'publickey': {'type': 'string'},
    'email': {'type': 'string'},
    'uuid': {'type': 'string'}
  }
}
```

Defined objects can also be embedded into other objects by using the *objectid* tag. This is later demonstrated between the profile and the user objects (see Objects 4.1 and 4.2).

When using the objects, it is assumed that one can implement the typical CRUD actions using HTTP methods mapped as follows:

GET	profile	Retrieves a list of profile
GET	profile12	Retrieves a specific profile
POST	profile	Creates a new profile
PUT	profile12	Updates profile #12
PATCH	profile12	Partially updates profile #12
DELETE	profile12	Deletes profile #12

In the reference implementation in this document, these methods are provided automatically.

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>



## 3.5 INTERFACE COMPLIANCY

Due to the easy extensibility of the objects in this document and their implicit interfaces, it is important to introduce a terminology that allows the definition of interface compliancy. The Subgroup defines three levels of interface compliance as follows:

- **Full Compliance:** These are reference implementations that provide full compliance to the objects defined in this document. A version number will be added to assure that the snapshot in time of the objects is associated with the version. This reference implementation will implement all objects.
- **Partial Compliance:** These are reference implementations that provide partial compliance to the objects defined in this document. A version number will be added to assure that the snapshot in time of the objects is associated with the version. This reference implementation will implement a partial list of the objects. A document will be generated during the reference implementation that lists all objects defined, but also lists the objects that are not defined by the reference architecture. The document will outline which objects and interfaces have been implemented.
- **Full and Extended Compliance:** These are interfaces that in addition to the full compliance also introduce additional interfaces and extend them. A document will be generated during the reference implementation that lists the differences to the document defined here.

The documents generated during the reference implementation can then be forwarded to the Reference Architecture Subgroup for further discussion and for possible future modifications based on additional practical user feedback.

# 4 SPECIFICATION

Several objects defined in this document are used across the NBDRA. Therefore, the objects have not been organized by NBDRA components, as shown in Figure 1, but rather grouped by functional use as summarized in Figure 2.

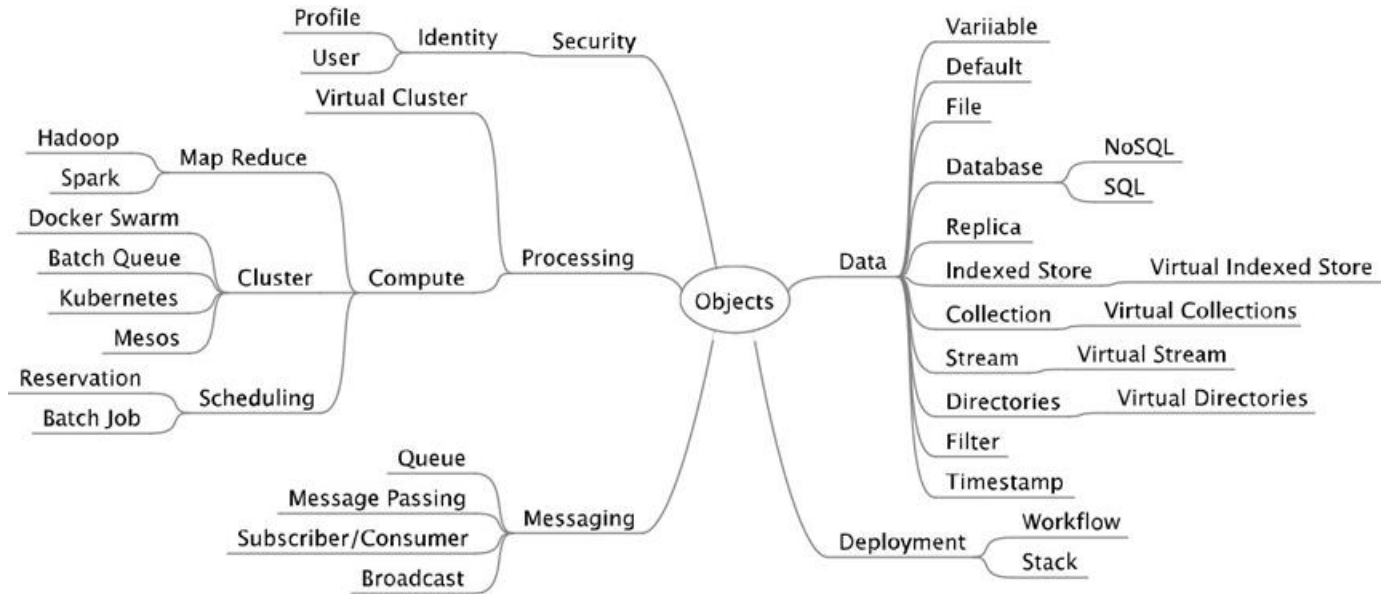


Figure 2: NIST Big Data Reference Architecture Interfaces

## 4.1 IDENTITY

In a multiuser environment, a simple mechanism is used in this document for associating objects and data to a particular person or group. While these efforts do not intend to replace more elaborate solutions such as proposed by eduPerson [12] or others, a very simple way was chosen to distinguish users. Therefore, the following sections introduce a number of simple objects including a profile and a user.

### 4.1.1 PROFILE

A profile defines the identity of an individual. It contains name and email information. It may have an optional unique user identification (uuid) and/or use a unique email to distinguish a user. Profiles are used to identify different users.

#### Object 4.1: Profile

```

1  {
2    "profile": {
3      "description": "The Profile of a user",
4      "uuid": "jshdjkdh...",
5      "context": "resource",
6      "email": "laszewski@gmail.com",
7      "firstname": "Gregor",
8      "lastname": "von Laszewski",
9      "username": "gregor",
10     "publickey": "ssh ...."

```

```
11 }
12 }
```

### 4.1.2 USER

In contrast to the profile, a user contains additional attributes that define the role of the user within the multiuser system. The user associates different roles to individuals. These roles potentially have gradations of responsibility and privilege.

#### Object 4.2: User

```
{
  "user": {
    "profile": "objectid:profile",
    "roles": ["admin"]
  }
}
```

### 4.1.3 ORGANIZATION

An important concept in many applications is the management of a group of users in an organization that manages a Big Data application or infrastructure. User group management can be achieved through three concepts. First, it can be achieved by using the profile and user resources itself as they contain the ability to manage multiple users as part of the REST interface. The second concept is to create a (virtual) organization that lists all users within the virtual organization. The third concept is to introduce groups and roles either as part of the user definition or as part of a simple list similar to the organization.

#### Object 4.3: Organization

```
{
  "organization": {
    "users": [
      "objectid:user"
    ]
  }
}
```

The profile, user, and organization concepts allow for the clear definition of various roles such as data provider, data consumer, data curator, and others. These concepts also allow for the creation of services that restrict data access by role or organizational affiliation.

### 4.1.4 GROUP/ROLE

A group contains a number of users. It is used to manage authorized services.

#### Object 4.4: Group

```
{
  "group": {
    "name": "users",
    "description": "This group contains all users",
    "users": [
      "objectid:user"
    ]
  }
}
```

A role is a further refinement of a group. Group members can have specific roles. For example, a group of users can be assigned a role that allows access to a repository. More specifically, the role would define a user's read and write privileges to the data within the repository.

#### Object 4.5: Role

```
{
  "role": {
    "name": "editor",
    "description": "This role contains all editors",
    "users": [
      "objectid:user"
    ]
  }
}
```

## 4.2 DATA

Data for Big Data applications are delivered through data providers. They can be either local data providers, data contributed by a user, or distributed data providers, data on the Internet. Currently, the focus is on an elementary set of abstractions related to data providers that offer the ability to utilize variables, files, virtual data directories, data streams, and data filters.

- **Variables** are used to hold specific contents that are associated in programming language as variables. A variable has a name, value, and type.
- **Defaults** are a special type of variables that allow adding of a context. Defaults can be created for different contexts.
- **Files** are used to represent information collected within the context of classical files in an operating system.
- **Directories** are locations for storing and organizing multiple files on a compute resource.
- **Virtual Directories** are collections of endpoints to files. Files in a virtual directory may be located on different resources. For this initial purpose, the distinction between virtual and non-virtual directories is nonessential, and the focus will be on abstracting all directories to be virtual. Therefore, the files could be physically hosted on different disks. However, it is important to note that virtual data directories can hold more than files; they can also contain data streams and data filters.
- **Streams** are services that offer the consumer a flow of data. Streams may allow the initiation of filters to reduce the amount of data requested by the consumer. Stream filters operate in streams or on files converting them to streams.
- **Batch Filters** operate on streams and on files, working in the background and delivering files as output. In contrast to streams, batch filters process the dataset and return an output after all operations have been applied.
- **Indexed Stores** are storage systems for objects that can be accessed by an index for each object. Search and filter functions are integrated into indexed stores to allow identification of objects.
- **Databases** refer to traditional databases but also to NoSQL.
- **Collections** are an agglomeration of any type of data.
- **Replicas** are duplication of data objects created to avoid overhead due to network or other physical restrictions on a remote resource.

### 4.2.1 TIMESTAMP

Often data needs to be timestamped to indicate when it has been accessed, created, or modified. All objects defined in this document will have, in their final version, a timestamp.

Object 4.6: TimeStamp
<pre> {   "timestamp": {     "accessed": "1.1.2017:05:00:00:EST",     "created": "1.1.2017:05:00:00:EST",     "modified": "1.1.2017:05:00:00:EST"   } }                     </pre>

## 4.2.2 VARIABLES

Variables are used to store simple values. Each variable can have a type, which is also provided as demonstrated in the object below. The variable value format is defined as string to allow maximal probability.

Object 4.7: Variables
<pre> {   "var": {     "name": "name of the variable",     "value": "the value of the variable as string",     "type": "the datatype of the variable such as int, str, float, ..."   } }                     </pre>

## 4.2.3 DEFAULT

A default is a special variable that has a context associated with it. This allows one to define values that can be easily retrieved based on the associated context. For example, a default could be the image name for a cloud where the context is defined by the cloud name.

Object 4.8: Default
<pre> {   "default": {     "value": "string",     "name": "string",     "context": "string - defines the context of the default (user, cloud, ...)"   } }                     </pre>

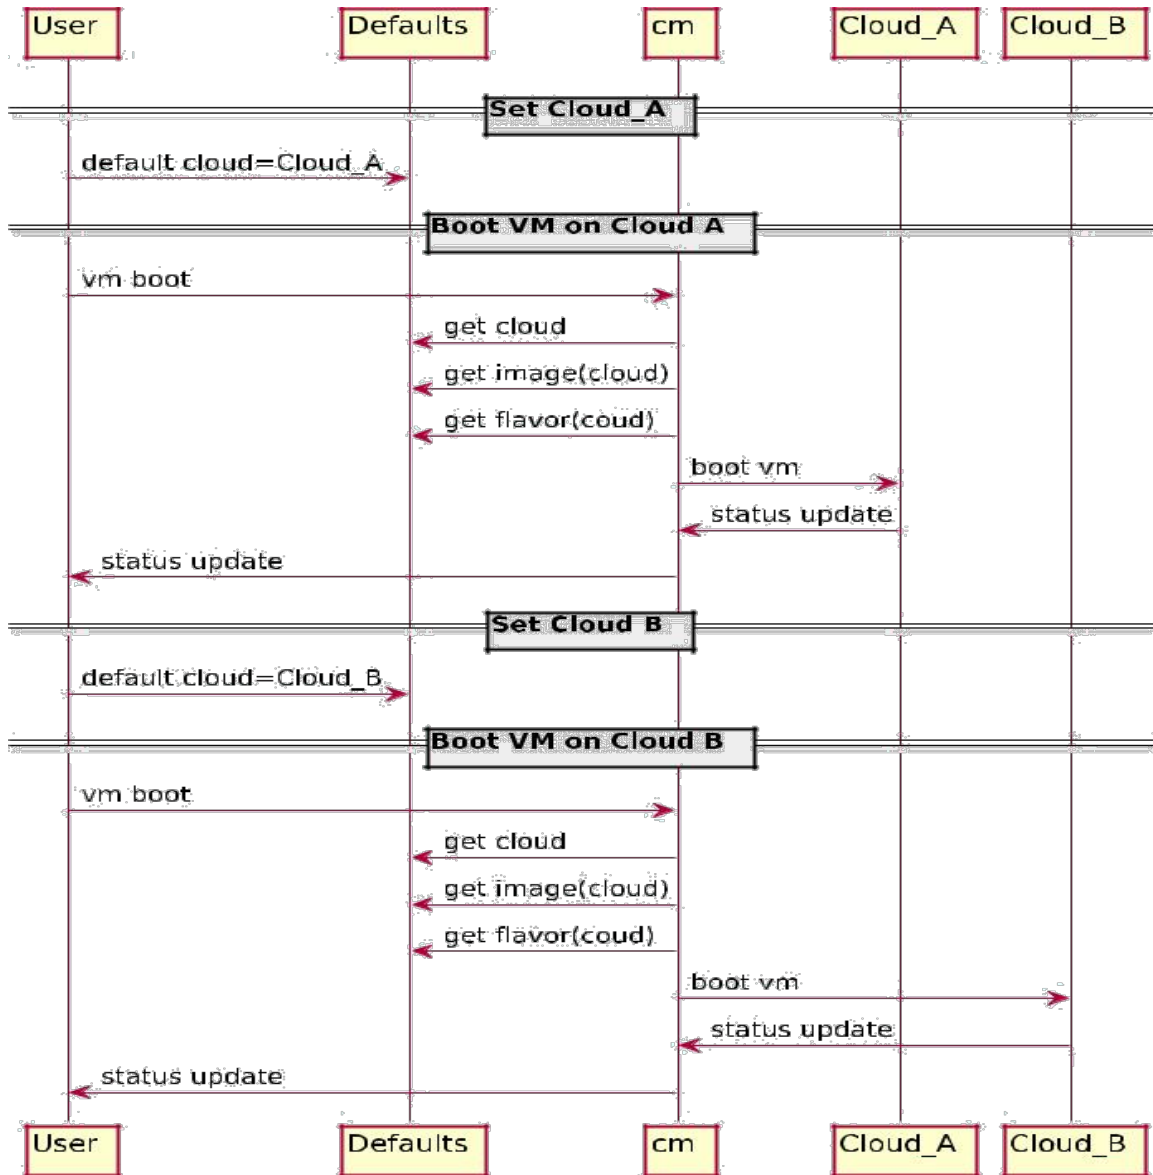


Figure 3: Booting a VM from Defaults

#### 4.2.4 FILE

A file is a computer resource allowing storage of data that is being processed. The interface to a file provides the mechanism to appropriately locate a file in a distributed system. File identification includes the name, endpoint, checksum, and size. Additional parameters, such as the last access time, could also be stored. The interface only describes the location of the file.

The *file* object has *name*, *endpoint* (location), *size* in GB, MB, Byte, *checksum* for integrity check, and last *accessed* timestamp.

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object 4.9: File**

```

{
  "file": {
    "name": "report.dat",
    "endpoint": "file://gregor@machine.edu:/data/report.dat",
    "checksum": {"sha256": "c01b39c7a35ccc .....
ebfeb45c69f08e17dfe3ef375a7b"},
    "accessed": "1.1.2017:05:00:00:EST",
    "created": "1.1.2017:05:00:00:EST",
    "modified": "1.1.2017:05:00:00:EST",
    "size": ["GB", "Byte"]
  }
}
    
```

**4.2.5 ALIAS**

A data object could have one alias or even multiple ones. The reason for an alias is that a file may have a complex name but a user may want to refer to that file in a name space that is suitable for the user's application.

**Object 4.10: File Alias**

```

{
  "alias": {
    "name": "a better name for the object",
    "origin": "the original object name"
  }
}
    
```

**4.2.6 REPLICA**

In many distributed systems, it is important that a file can be replicated among different systems to provide faster access. It is important to provide a mechanism to trace the pedigree of the file while pointing to its original source. A replica can be applied to all data types introduced in this document.

**Object 4.11: Replica**

```

{
  "replica": {
    "name": "replica_report.dat",
    "replica": "report.dat",
    "endpoint": "file://gregor@machine.edu:/data/replica_report.dat",
    "checksum": {
      "md5": "8c324f12047dc2254b74031b8f029ad0"
    },
    "accessed": "1.1.2017:05:00:00:EST",
    "size": [
      "GB",
      "Byte"
    ]
  }
}
    
```

### 4.2.7 VIRTUAL DIRECTORY

A virtual directory is a collection of files or replicas of the files. A virtual directory can contain a number of entities including files, streams, and other virtual directories as part of a collection. The element in the collection can either be defined by uuid or by name.

**Object 4.12: Virtual Directory**

```
{
  "virtual_directory": {
    "name": "data",
    "endpoint": "http://.../data/",
    "protocol": "http",
    "collection": [
      "report.dat",
      "file2"
    ]
  }
}
```

### 4.2.8 DATABASE

A *database* could have a name, an *endpoint* (e.g., host, port), and a protocol used (e.g., SQL, mongo).

**Object 4.13: Database**

```
{
  "database": {
    "name": "data",
    "endpoint": "http://.../data/",
    "protocol": "mongo"
  }
}
```

### 4.2.9 STREAM

The stream object describes a data flow, providing information about the rate and number of items exchanged while issuing requests to the stream. A stream may return data items in a specific format that is defined by the stream.

**Object 4.14: Stream**

```
{
  "stream": {
    "name": "name of the variable",
    "format": "the format of the data exchanged in the stream",
    "attributes": {
      "rate": 10,
      "limit": 1000
    }
  }
}
```

### 4.2.10 FILTER

Filters can operate on a variety of objects and reduce the information received based on a search criterion.

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>



**Object 4.15: Filter**

```

{
  "filter": {
    "name": "name of the filter",
    "function": "the function of the data exchanged in the stream"
  }
}
    
```

### 4.3 VIRTUAL CLUSTER

One of the essential features for Big Data is the creation of a Big Data analysis cluster. A virtual cluster combines resources that generally are used to serve the Big Data application and can constitute a variety of data analysis nodes that together build the virtual cluster. Instead of focusing only on the deployment of a physical cluster, the creation of a virtual cluster can be instantiated on a number of different platforms. Such platforms include clouds, containers, physical hardware, or a mix thereof to support different aspects of the Big Data application.

Figure 4 illustrates the process for allocating and provisioning a virtual cluster. The user defines the desired physical properties of the cluster (e.g., CPU, memory, disk) and the intended configuration (e.g., software, users). After requesting the stack to be deployed, cloudmesh allocates the machines by matching the desired properties with the available images and booting. The stack definition is then parsed and then evaluated to provision the cluster.

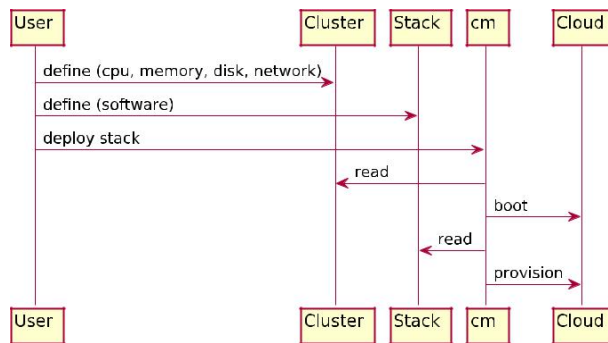


Figure 4: Allocating and Provisioning a Virtual Cluster

#### 4.3.1 VIRTUAL CLUSTER

A virtual cluster is an agglomeration of virtual compute nodes that constitute the cluster. Nodes can be assembled to be bare metal, VMs, and containers. A virtual cluster contains a number of virtual compute nodes.

The virtual cluster object has name, label, endpoint, and provider. The *endpoint* defines a mechanism to connect to it. The *provider* defines the nature of the cluster (e.g., it is a virtual cluster on an OpenStack cloud, or from AWS, or a bare metal cluster).

To manage the cluster, it can have a frontend node that is used to manage other nodes. Authorized keys within the definition of the cluster allow administrative functions, while authorized keys on a compute node allow login and use functionality of the virtual nodes.

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object 4.16: Virtual Cluster**

```
{
  "virtual_cluster": {
    "name": "myvirtualcluster",
    "label": "C0",
    "uuid": "sgdlsjlaj...",
    "endpoint": {
      "passwd": "secret",
      "url": "https:..."
    },
    "provider": "virtual_cluster_provider:openstack",
    "frontend": "objectid:virtual_machine",
    "authorized_keys": ["objectid:sshkey"],
    "nodes": [
      "objectid:virtual_machine"
    ]
  }
}
```

**Object 4.17: Virtual Cluster Provider**

```
"virtual_cluster_provider": "aws" | "azure" | "google" | "comet" | "openstack"
```

**4.3.2 COMPUTE NODE**

Compute nodes are used to conduct compute and data functions. They are of a specific *kind*. For example, compute nodes could be a virtual machine (VM), bare metal, or part of a predefined virtual cluster framework.

Compute nodes are a representation of a computer system (physical or virtual). A very basic set of information about the compute node is maintained in this document. It is expected that, through the endpoint, the VM can be introspected and more detailed information can be retrieved. A compute node has name, label, a flavor, network interface cards (NICs), and other relevant information.

**Object 4.18: Compute Node of a Virtual Cluster**

```
{
  "compute_node": {
    "name": "vm1",
    "label": "gregor-vm001",
    "uuid": "sgklfgslakj...",
    "kind": "vm",
    "flavor": ["objectid:flavor"],
    "image": "Ubuntu-16.04",
    "secgroups": ["objectid:secgroup"],
    "nics": ["objectid:nic"],
    "status": "",
    "loginuser": "ubuntu",
    "status": "active",
    "authorized_keys": ["objectid:sshkey"],
    "metadata": {
      "owner": "gregor",
      "experiment": "exp-001"
    }
  }
}
```

```
}

```

### 4.3.3 FLAVOR

The flavor specifies elementary information about the compute node, such as memory and number of cores, as well as other attributes that can be added. Flavors are essential to size a virtual cluster appropriately.

#### Object 4.19: Flavor

```
{
  "flavor": {
    "name": "flavor1",
    "label": "2-4G-40G",
    "uuid": "sgklfgslakj....",
    "ncpu": 2,
    "ram": "4G",
    "disk": "40G"
  }
}
```

### 4.3.4 NETWORK INTERFACE CARD

To interact between the nodes, a network interface is needed. Such a network interface, specified on a virtual machine with a NIC object, is showcased in Object 4.20.

#### Object 4.20: Network Interface Card

```
{
  "nic": {
    "name": "eth0",
    "type": "ethernet",
    "mac": "00:00:00:11:22:33",
    "ip": "123.123.1.2",
    "mask": "255.255.255.0",
    "broadcast": "123.123.1.255",
    "gateway": "123.123.1.1",
    "mtu": 1500,
    "bandwidth": "10Gbps"
  }
}
```

### 4.3.5 KEY

Many services and frameworks use Secure Shell (SSH) keys to authenticate. To allow the convenient storage of the public key, the sshkey object can be used (see Object 4.21).

#### Object 4.21: Key

```
{
  "sshkey": {
    "comment": "string",
    "source": "string",
    "uri": "string",
    "value": "ssh-rsa AAA.....",
    "fingerprint": "string, unique"
  }
}
```

### 4.3.6 SECURITY GROUPS

To allow secure communication between the nodes, security groups are introduced. They define the typical security groups that will be deployed once a compute node is specified. The security group object is depicted in Object 4.22.

Object 4.22: Security Groups
<pre> {   "secgroup":   {     "ingress": "0.0.0.0/32",     "egress": "0.0.0.0/32",     "ports": 22,     "protocols": "tcp"   } }                     </pre>

## 4.4 IAAS

Although Section 4.3 defines a general virtual cluster useful for Big Data, sometimes the need exists to specifically utilize Infrastructure as a Service (IaaS) frameworks, such as OpenStack, Azure, and others. To do so, it is beneficial to be able to define virtual clusters using these frameworks. Hence, this subsection defines interfaces related to IaaS frameworks. This includes specific objects useful for OpenStack, Azure, and AWS, as well as others. The definition of the objects used between the clouds to manage them, are different and not standardized. In this case, the objects support functions such as starting, stopping, suspending resuming, migration, network configuration, assigning of resources, assigning of operating systems for and others for the VMs.

Inspecting other examples, such as *LibCloud*, shows the definition of generalized objects are discovered, which are augmented with extra fields to specifically integrate with the various frameworks. When working with cloudmesh, it is sufficient to be able to specify a cloud based on a cloud-specific action. Actions include boot, terminate, suspend, resume, assign network intrusion prevention system, and add users.

To support such actions, objects can be selected based on the IaaS type in use when invoked. The following subsections list these objects as used in LibCloud, OpenStack, and Azure.

### 4.4.1 LIBCLOUD

Libcloud is a Python library for interacting with different cloud service providers. It uses a unified API that exposes similar access to a variety of clouds. Internally, it uses objects that can interface with different IaaS frameworks. However, as these frameworks are different from each other, specific adaptations are done for each IaaS, mostly reflected in the LibCloud Node (see Section 4.4.1.5).

#### 4.4.1.1 Challenges

For time considerations, LibCloud was used for some time practically in various versions of cloudmesh. However, it became apparent that at times, the representation and functionality provided by LibCloud, for reference implementations, did not support some advanced aspects provided by the native cloud objects. Depending on the application, libraries for interfacing with different frameworks, direct utilization of the native objects, and interfaces provided by a particular IaaS framework could all be viable options. Additional interfaces have been introduced in Sections 4.4.2 and 4.4.3. Additional sections addressing other IaaS frameworks may be integrated in the future.

#### 4.4.1.2 *LibCloud Flavor*

The object referring to flavors is listed in Object 4.23.

##### Object 4.23: Libcloud Flavor

```
{
  "libcloud_flavor": {
    "bandwidth": "string",
    "disk": "string",
    "uuid": "string",
    "price": "string",
    "ram": "string",
    "cpu": "string",
    "flavor_id": "string"
  }
}
```

#### 4.4.1.3 *LibCloud Image*

The object referring to images is listed in Object 4.24.

##### Object 4.24: Libcloud Image

```
{
  "libcloud_image": {
    "username": "string",
    "status": "string",
    "updated": "string",
    "description": "string",
    "owner_alias": "string",
    "kernel_id": "string",
    "ramdisk_id": "string",
    "image_id": "string",
    "is_public": "string",
    "image_location": "string",
    "uuid": "string",
    "created": "string",
    "image_type": "string",
    "hypervisor": "string",
    "platform": "string",
    "state": "string",
    "architecture": "string",
    "virtualization_type": "string",
    "owner_id": "string"
  }
}
```

#### 4.4.1.4 *LibCloud VM*

The object referring to virtual machines is listed in the Object 4.25.

**Object 4.25: LibCloud VM**

```

{
  "libcloud_vm": {
    "username": "string",
    "status": "string",
    "root_device_type": "string",
    "image": "string",
    "image_name": "string",
    "image_id": "string",
    "key": "string",
    "flavor": "string",
    "availability": "string",
    "private_ips": "string",
    "group": "string",
    "uuid": "string",
    "public_ips": "string",
    "instance_id": "string",
    "instance_type": "string",
    "state": "string",
    "root_device_name": "string",
    "private_dns": "string"
  }
}

```

**4.4.1.5 LibCloud Node**

Virtual machines for the various clouds have additional attributes that are summarized in Object 4.25. These attributes will be integrated into the VM object in the future.

**Object 4.26: LibCloud Node**

```

{
  "LibCloudNode": {
    "id": "instance_id",
    "name": "name",
    "state": "state",
    "public_ips": ["111.222.111.1"],
    "private_ips": ["192.168.1.101"],
    "driver": "connection.driver",
    "created_at": "created_timestamp",
    "extra": {
    }
  },
  "ec2NodeExtra": {
    "block_device_mapping": "deviceMapping",
    "groups": ["security_group1", "security_group2"],
    "network_interfaces": ["nic1", "nic2"],
    "product_codes": "product_codes",
    "tags": ["tag1", "tag2"]
  },
  "OpenStackNodeExtra": {
    "addresses": ["addresses"],
    "hostId": "hostId",
    "access_ip": "accessIPv4",
    "access_ipv6": "accessIPv6",
  }
}

```

**Object 4.26: LibCloud Node**

```

    "tenantId": "tenant_id",
    "userId": "user_id",
    "imageId": "image_id",
    "flavorId": "flavor_id",
    "uri": "",
    "service_name": "",
    "metadata": ["metadata"],
    "password": "adminPass",
    "created": "created",
    "updated": "updated",
    "key_name": "key_name",
    "disk_config": "diskConfig",
    "config_drive": "config_drive",
    "availability_zone": "availability_zone",
    "volumes_attached": "volumes_attached",
    "task_state": "task_state",
    "vm_state": "vm_state",
    "power_state": "power_state",
    "progress": "progress",
    "fault": "fault"
  },
  "AzureNodeExtra": {
    "instance_endpoints": "instance_endpoints",
    "remote_desktop_port": "remote_desktop_port",
    "ssh_port": "ssh_port",
    "power_state": "power_state",
    "instance_size": "instance_size",
    "ex_cloud_service_name": "ex_cloud_service_name"
  },
  "GCENodeExtra": {
    "status": "status",
    "statusMessage": "statusMessage",
    "description": "description",
    "zone": "zone",
    "image": "image",
    "machineType": "machineType",
    "disks": "disks",
    "networkInterfaces": "networkInterfaces",
    "id": "node_id",
    "selfLink": "selfLink",
    "kind": "kind",
    "creationTimestamp": "creationTimestamp",
    "name": "name",
    "metadata": "metadata",
    "tags_fingerprint": "fingerprint",
    "scheduling": "scheduling",
    "deprecated": "True or False",
    "canIpForward": "canIpForward",
    "serviceAccounts": "serviceAccounts",
    "boot_disk": "disk"
  }
}

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

## 4.4.2 OPENSTACK

Objects related to OpenStack VMs are summarized in this section.

### 4.4.2.1 OpenStack Flavor

The object referring to flavors is listed in Object 4.27.

#### Object 4.27: OpenStack Flavor

```
{
  "openstack_flavor": {
    "os_flv_disabled": "string",
    "uuid": "string",
    "os_flv_ext_data": "string",
    "ram": "string",
    "os_flavor_acces": "string",
    "vcpus": "string",
    "swap": "string",
    "rxtx_factor": "string",
    "disk": "string"
  }
}
```

### 4.4.2.2 OpenStack Image

The object referring to images is listed in Object 4.28.

#### Object 4.28: OpenStack Image

```
{
  "openstack_image": {
    "status": "string",
    "username": "string",
    "updated": "string",
    "uuid": "string",
    "created": "string",
    "minDisk": "string",
    "progress": "string",
    "minRam": "string",
    "os_image_size": "string",
    "metadata": {
      "image_location": "string",
      "image_state": "string",
      "description": "string",
      "kernel_id": "string",
      "instance_type_id": "string",
      "ramdisk_id": "string",
      "instance_type_name": "string",
      "instance_type_rxtx_factor": "string",
      "instance_type_vcpus": "string",
      "user_id": "string",
      "base_image_ref": "string",
      "instance_uuid": "string",
      "instance_type_memory_mb": "string",
      "instance_type_swap": "string",
      "image_type": "string",
    }
  }
}
```



This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

```

Object 4.28: OpenStack Image
    "instance_type_ephemeral_gb": "string",
    "instance_type_root_gb": "string",
    "network_allocated": "string",
    "instance_type_flavorid": "string",
    "owner_id": "string"
  }
}
}

```

**4.4.2.3 OpenStack VM**

The object referring to VMs is listed in Object 4.29.

```

Object 4.29: OpenStack VM
{
  "openstack_vm": {
    "username": "string",
    "vm_state": "string",
    "updated": "string",
    "hostId": "string",
    "availability_zone": "string",
    "terminated_at": "string",
    "image": "string",
    "floating_ip": "string",
    "diskConfig": "string",
    "key": "string",
    "flavor_id": "string",
    "user_id": "string",
    "flavor": "string",
    "static_ip": "string",
    "security_groups": "string",
    "volumes_attached": "string",
    "task_state": "string",
    "group": "string",
    "uuid": "string",
    "created": "string",
    "tenant_id": "string",
    "accessIPv4": "string",
    "accessIPv6": "string",
    "status": "string",
    "power_state": "string",
    "progress": "string",
    "image_id": "string",
    "launched_at": "string",
    "config_drive": "string"
  }
}

```

**4.4.3 AZURE**

Objects related to Azure virtual machines are summarized in this section.

### 4.4.3.1 Azure Size

The object referring to the image size machines is listed in Object 4.30.

#### Object 4.30: Azure-Size

```
{
  "azure-size": {
    "_uuid": "None",
    "name": "D14 Faster Compute Instance",
    "extra": {
      "cores": 16,
      "max_data_disks": 32
    },
    "price": 1.6261,
    "ram": 114688,
    "driver": "libcloud",
    "bandwidth": "None",
    "disk": 127,
    "id": "Standard_D14"
  }
}
```

### 4.4.3.2 Azure Image

The object referring to the images machines is listed in Object 4.31.

#### Object 4.31: Azure-Image

```
{
  "azure_image": {
    "_uuid": "None",
    "driver": "libcloud",
    "extra": {
      "affinity_group": "",
      "category": "Public",
      "description": "Linux VM image with coreclr-x64-beta5-11624 installed to /opt/dnx. This image is based on Ubuntu 14.04 LTS, with prerequisites of CoreCLR installed. It also contains PartsUnlimited demo app which runs on the installed coreclr. The demo app is installed to /opt/demo. To run the demo, please type the command /opt/demo/Kestrel in a terminal window. The website is listening on port 5004. Please enable or map an endpoint of HTTP port 5004 for your azure VM.",
      "location": "East Asia;Southeast Asia;Australia East;Australia Southeast;Brazil South;North Europe;West Europe;Japan East;Japan West;Central US;East US 2; North Central US;South Central US;West US",
      "media_link": "",
      "os": "Linux",
      "vm_image": "False"
    },
    "id": "03f55de797f546a1b29d1....",
    "name": "CoreCLR x64 Beta5 (11624) with PartsUnlimited Demo App on Ubuntu Server 14.04 LTS"
  }
}
```

### 4.4.3.3 Azure VM

The object referring to the virtual machines is listed in Object 4.32.

**Object 4.32: Azure VM**

```
{
  "azure-vm": {
    "username": "string",
    "status": "string",
    "deployment_slot": "string",
    "cloud_service": "string",
    "image": "string",
    "floating_ip": "string",
    "image_name": "string",
    "key": "string",
    "flavor": "string",
    "resource_location": "string",
    "disk_name": "string",
    "private_ips": "string",
    "group": "string",
    "uuid": "string",
    "dns_name": "string",
    "instance_size": "string",
    "instance_name": "string",
    "public_ips": "string",
    "media_link": "string"
  }
}
```

## 4.5 COMPUTE SERVICES

### 4.5.1 BATCH QUEUE

Computing jobs that can run without end user interaction, or are scheduled based on resource permission, are called batch jobs. Batch jobs are used to minimize human interaction and allow the submission and scheduling of many jobs in parallel while attempting to utilize the resources through a resource scheduler more efficiently or simply in sequential order. Batch processing is not to be underestimated even in today’s shifting Internet of Things environment towards clouds and containers. This is based on the fact that for some applications, resources managed by batch queues are highly optimized and in many cases, provide significant performance advantages. Disadvantages include the limited and preinstalled software stacks that, in some cases, do not allow the latest applications to run.

**Object 4.33: Batch Job**

```
{
  "batchjob": {
    "output_file": "string",
    "group": "string",
    "job_id": "string",
    "script": "string, the batch job script",
    "cmd": "string, executes the cmd, if None path is used",
    "queue": "string",
    "cluster": "string",
    "time": "string",
    "path": "string, path of the batchjob, if non cmd is used",
  }
}
```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

```

        "nodes": "string",
        "dir": "string"
    }
}
    
```

### 4.5.2 RESERVATION

Some services may consume a considerable amount of resources, necessitating the reservation of resources. For this purpose, a reservation object (Object 4.34) has been introduced.

#### Object 4.34: Reservation

```

{
  "reservation": {
    "service": "name of the service",
    "description": "what is this reservation for",
    "start_time": ["date", "time"],
    "end_time": ["date", "time"]
  }
}
    
```

## 4.6 CONTAINERS

The following defines the container object.

#### Object 4.35: Container

```

{
  "container": {
    "name": "container1",
    "endpoint": "http://.../container/",
    "ip": "127.0.0.1",
    "label": "server-001",
    "memoryGB": 16
  }
}
    
```

## 4.7 DEPLOYMENT

A *deployment* consists of the resource *cluster*, the location *provider* (e.g., OpenStack), and software *stack* to be deployed (e.g., Hadoop, Spark).

#### Object 4.36: Deployment

```

{
  "deployment": {
    "cluster": [{ "name": "myCluster"},
                { "id" : "cm-0001"}
    ],
    "stack": {
      "layers": [
        "zookeeper",
        "hadoop",
        "spark",
        "postgresql"
      ]
    }
  }
}
    
```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

```

Object 4.36: Deployment
    "parameters": {
      "hadoop": { "zookeeper.quorum": [ "IP", "IP", "IP"]
    }
  }
}
    
```

## 4.8 MAP/REDUCE

The *Map/Reduce* deployment has as inputs parameters defining the applied function and the input data. Both function and data objects define a *source* parameter, which specifies the location from which it is retrieved. For instance, the `file://` Uniform Resource Identifier (URI) indicates sending a directory structure from the local file system, and the `ftp://` indicates that the data should be fetched from a File Transfer Protocol (FTP) resource. It is the framework's responsibility to materialize an instantiation of the desired environment along with the function and data.

```

Object 4.37: Map/Reduce
    {
      "mapreduce": {
        "function": {
          "source": "file://.",
          "args": {}
        },
        "data": {
          "source": "ftp:///...",
          "dest": "/data"
        },
        "fault_tolerant": true,
        "backend": {"type": "hadoop"}
      }
    }
    
```

Additional parameters include the *fault\_tolerant* and *backend* parameters. The former flag indicates if the *Map/Reduce* deployment should operate in a fault tolerant mode. For instance, in the case of Hadoop, this may mean configuring automatic failover of name nodes using Zookeeper. The *backend* parameter accepts an object describing the system providing the *Map/Reduce* workflow. This may be a native deployment of Hadoop, or a special instantiation using other frameworks such as Mesos.

A function prototype is defined earlier. Key properties are that functions describe their input parameters and generated results. For input parameters, the *buildInputs* and *systemBuildInputs* respectively describe the objects that should be evaluated and system packages that should be present before this function can be installed. The *eval* attribute describes how to apply this function to its input data. Parameters affecting the evaluation of the function may be passed in as the *args* attribute. The results of the function application can be accessed via the *outputs* object, which is a mapping from arbitrary keys (e.g., data, processed, model) to an object representing the result.

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object 4.38: Map/Reduce Function**

```
{
  "mapreduce_function": {
    "name": "name of this function",
    "description": "These should be self-describing",
    "source": "a URI to obtain the resource",
    "install": {
      "description": "instructions to install the source if needed",
      "script": "source://install.sh"
    },
    "eval": {
      "description": "How to evaluate this function",
      "script": "source://run.sh"
    },
    "args": [
      {
        "argument": "value"
      }
    ],
    "buildInputs": [
      "list of dependent objects"
    ],
    "systemBuildInputs": [
      "list of packages"
    ],
    "outputs": {
      "key": "value"
    }
  }
}
```

One example function is the *NoOp* function shown in Object 4.39. In the case of undefined arguments, the parameters default to an identity element. In the case of mappings, the identity element is the empty mapping while for lists, the identity element is the empty list.

**Object 4.39: Map/Reduce NoOp**

```
{
  "mapreduce_noop": {
    "name": "noop",
    "description": "A function with no effect"
  }
}
```

**4.8.1 HADOOP**

A *Hadoop* definition defines which *deployer* to use, the *parameters* of the deployment, and the system packages *required*. For each requirement, it could have attributes such as the library origin, version, and other attributes (see Object 4.40).

**Object 4.40: Hadoop**

```

{
  "hadoop": {
    "deployers": {
      "ansible": "git://github.com/cloudmesh_roles/hadoop"
    },
    "requires": {
      "java": {
        "implementation": "OpenJDK",
        "version": "1.8",
        "zookeeper": "TBD",
        "supervisord": "TBD"
      }
    },
    "parameters": {
      "num_resourcemangers": 1,
      "num_namenodes": 1,
      "use_yarn": false,
      "use_hdfs": true,
      "num_datanodes": 1,
      "num_historyservers": 1,
      "num_journalnodes": 1
    }
  }
}

```

**4.9 MICROSERVICE**

As part of microservices, a function with parameters that can be invoked has been defined. To describe such services, the Object 4.41 was created. Defining multiple services facilitates the finding of the microservices and the use as part of a microservice-based implementation.

**Object 4.41: Microservice**

```

{
  "microservice" :{
    "name": "ms1",
    "endpoint": "http://.../ms/",
    "function": "microservice spec"
  }
}

```

# 5 STATUS CODES AND ERROR RESPONSES

---

In case of an error or a successful response, the response header contains a HTTP code (see <https://tools.ietf.org/html/rfc7231>). The response body usually contains the following:

- The HTTP response code;
- An accompanying message for the HTTP response code; and
- A field or object where the error occurred.

*Table 1: HTTP Response Codes*

<b>HTTP</b>	<b>Response Description Code</b>
200	<i>OK</i> success code, for GET or HEAD request.
201	<i>Created</i> success code, for POST request.
204	<i>No Content</i> success code, for DELETE request.
300	The value returned when an external ID exists in more than one record.
304	The request content has not changed since a specified date and time.
400	The request could not be understood.
401	The session ID or OAuth token used has expired or is invalid.
403	The request has been refused.
404	The requested resource could not be found.
405	The method specified in the Request-Line isn't allowed for the resource specified in the URI.
415	The entity in the request is in a format that's not supported by the specified method.



# Appendix A: Schema

Object A.1 showcases the schema generated from the objects defined in this document.

## Object A.1: Schema

```

container = {
  'schema': {
    'ip': {
      'type': 'string'
    },
    'endpoint': {
      'type': 'string'
    },
    'name': {
      'type': 'string'
    },
    'memoryGB': {
      'type': 'integer'
    },
    'label': {
      'type': 'string'
    }
  }
}

stream = {
  'schema': {
    'attributes': {
      'type': 'dict',
      'schema': {
        'rate': {
          'type': 'integer'
        },
        'limit': {
          'type': 'integer'
        }
      }
    },
    'name': {
      'type': 'string'
    },
    'format': {
      'type': 'string'
    }
  }
}

azure_image = {
  'schema': {
    '_uuid': {
      'type': 'string'
    },
  },

```

**Object A.1: Schema**

```

'driver': {
  'type': 'string'
},
'id': {
  'type': 'string'
},
'name': {
  'type': 'string'
},
'extra': {
  'type': 'dict',
  'schema': {
    'category': {
      'type': 'string'
    },
    'description': {
      'type': 'string'
    },
    'vm_image': {
      'type': 'string'
    },
    'location': {
      'type': 'string'
    },
    'affinity_group': {
      'type': 'string'
    },
    'os': {
      'type': 'string'
    },
    'media_link': {
      'type': 'string'
    }
  }
}
}
}
}

deployment = {
  'schema': {
    'cluster': {
      'type': 'list',
      'schema': {
        'type': 'dict',
        'schema': {
          'id': {
            'type': 'string'
          }
        }
      }
    }
  }
},
'stack': {
  'type': 'dict',

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

'schema': {
  'layers': {
    'type': 'list',
    'schema': {
      'type': 'string'
    }
  },
  'parameters': {
    'type': 'dict',
    'schema': {
      'hadoop': {
        'type': 'dict',
        'schema': {
          'zookeeper.quorum': {
            'type': 'list',
            'schema': {
              'type': 'string'
            }
          }
        }
      }
    }
  }
}

azure_size = {
  'schema': {
    'ram': {
      'type': 'integer'
    },
    'name': {
      'type': 'string'
    },
    'extra': {
      'type': 'dict',
      'schema': {
        'cores': {
          'type': 'integer'
        },
        'max_data_disks': {
          'type': 'integer'
        }
      }
    },
    'price': {
      'type': 'float'
    },
    '_uuid': {
      'type': 'string'
    }
  }
}

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

        'driver': {
            'type': 'string'
        },
        'bandwidth': {
            'type': 'string'
        },
        'disk': {
            'type': 'integer'
        },
        'id': {
            'type': 'string'
        }
    }
}

cluster = {
    'schema': {
        'provider': {
            'type': 'list',
            'schema': {
                'type': 'string'
            }
        },
        'endpoint': {
            'type': 'dict',
            'schema': {
                'passwd': {
                    'type': 'string'
                },
                'url': {
                    'type': 'string'
                }
            }
        },
        'name': {
            'type': 'string'
        },
        'label': {
            'type': 'string'
        }
    }
}

computer = {
    'schema': {
        'ip': {
            'type': 'string'
        },
        'name': {
            'type': 'string'
        },
        'memoryGB': {
            'type': 'integer'
        }
    }
}

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

    },
    'label': {
      'type': 'string'
    }
  }
}

mesos_docker = {
  'schema': {
    'container': {
      'type': 'dict',
      'schema': {
        'docker': {
          'type': 'dict',
          'schema': {
            'credential': {
              'type': 'dict',
              'schema': {
                'secret': {
                  'type': 'string'
                },
                'principal': {
                  'type': 'string'
                }
              }
            },
            'image': {
              'type': 'string'
            }
          }
        },
        'type': {
          'type': 'string'
        }
      }
    },
    'mem': {
      'type': 'float'
    },
    'args': {
      'type': 'list',
      'schema': {
        'type': 'string'
      }
    },
    'cpus': {
      'type': 'float'
    },
    'instances': {
      'type': 'integer'
    },
    'id': {
      'type': 'string'
    }
  }
}

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

    }
  }
}

file = {
  'schema': {
    'endpoint': {
      'type': 'string'
    },
    'name': {
      'type': 'string'
    },
    'created': {
      'type': 'string'
    },
    'checksum': {
      'type': 'dict',
      'schema': {
        'sha256': {
          'type': 'string'
        }
      }
    },
    'modified': {
      'type': 'string'
    },
    'accessed': {
      'type': 'string'
    },
    'size': {
      'type': 'list',
      'schema': {
        'type': 'string'
      }
    }
  }
}

reservation = {
  'schema': {
    'start_time': {
      'type': 'list',
      'schema': {
        'type': 'string'
      }
    },
    'description': {
      'type': 'string'
    },
    'service': {
      'type': 'string'
    },
    'end_time': {

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

        'type': 'list',
        'schema': {
            'type': 'string'
        }
    }
}

microservice = {
    'schema': {
        'function': {
            'type': 'string'
        },
        'endpoint': {
            'type': 'string'
        },
        'name': {
            'type': 'string'
        }
    }
}

flavor = {
    'schema': {
        'uuid': {
            'type': 'string'
        },
        'ram': {
            'type': 'string'
        },
        'label': {
            'type': 'string'
        },
        'ncpu': {
            'type': 'integer'
        },
        'disk': {
            'type': 'string'
        },
        'name': {
            'type': 'string'
        }
    }
}

virtual_directory = {
    'schema': {
        'endpoint': {
            'type': 'string'
        },
        'protocol': {
            'type': 'string'
        },
    },
}

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

        'name': {
            'type': 'string'
        },
        'collection': {
            'type': 'list',
            'schema': {
                'type': 'string'
            }
        }
    }
}

mapreduce_function = {
    'schema': {
        'name': {
            'type': 'string'
        },
        'outputs': {
            'type': 'dict',
            'schema': {
                'key': {
                    'type': 'string'
                }
            }
        },
        'args': {
            'type': 'list',
            'schema': {
                'type': 'dict',
                'schema': {
                    'argument': {
                        'type': 'string'
                    }
                }
            }
        },
        'systemBuildInputs': {
            'type': 'list',
            'schema': {
                'type': 'string'
            }
        },
        'source': {
            'type': 'string'
        },
        'install': {
            'type': 'dict',
            'schema': {
                'description': {
                    'type': 'string'
                },
                'script': {
                    'type': 'string'
                }
            }
        }
    }
}

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>



**Object A.1: Schema**

```

    }
  },
  'eval': {
    'type': 'dict',
    'schema': {
      'description': {
        'type': 'string'
      },
      'script': {
        'type': 'string'
      }
    }
  },
  'buildInputs': {
    'type': 'list',
    'schema': {
      'type': 'string'
    }
  },
  'description': {
    'type': 'string'
  }
}
}

virtual_cluster = {
  'schema': {
    'authorized_keys': {
      'type': 'list',
      'schema': {
        'type': 'objectid',
        'data_relation': {
          'resource': 'sshkey',
          'field': '_id',
          'embeddable': True
        }
      }
    },
    'endpoint': {
      'type': 'dict',
      'schema': {
        'passwd': {
          'type': 'string'
        },
        'url': {
          'type': 'string'
        }
      }
    },
    'frontend': {
      'type': 'objectid',
      'data_relation': {

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

        'resource': 'virtual_machine',
        'field': '_id',
        'embeddable': True
    }
},
'uuid': {
    'type': 'string'
},
'label': {
    'type': 'string'
},
'provider': {
    'type': 'string'
},
'nodes': {
    'type': 'list',
    'schema': {
        'type': 'objectid',
        'data_relation': {
            'resource': 'virtual_machine',
            'field': '_id',
            'embeddable': True
        }
    }
},
'name': {
    'type': 'string'
}
}
}

libcloud_flavor = {
    'schema': {
        'uuid': {
            'type': 'string'
        },
        'price': {
            'type': 'string'
        },
        'ram': {
            'type': 'string'
        },
        'bandwidth': {
            'type': 'string'
        },
        'flavor_id': {
            'type': 'string'
        },
        'disk': {
            'type': 'string'
        },
        'cpu': {
            'type': 'string'
        }
    }
}

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

    }
  }
}

LibCloudNode = {
  'schema': {
    'private_ips': {
      'type': 'list',
      'schema': {
        'type': 'string'
      }
    },
    'extra': {
      'type': 'dict',
      'schema': {}
    },
    'created_at': {
      'type': 'string'
    },
    'driver': {
      'type': 'string'
    },
    'state': {
      'type': 'string'
    },
    'public_ips': {
      'type': 'list',
      'schema': {
        'type': 'string'
      }
    },
    'id': {
      'type': 'string'
    },
    'name': {
      'type': 'string'
    }
  }
}

sshkey = {
  'schema': {
    'comment': {
      'type': 'string'
    },
    'source': {
      'type': 'string'
    },
    'uri': {
      'type': 'string'
    },
    'value': {
      'type': 'string'
    }
  }
}

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

        },
        'fingerprint': {
            'type': 'string'
        }
    }
}

timestamp = {
    'schema': {
        'accessed': {
            'type': 'string'
        },
        'modified': {
            'type': 'string'
        },
        'created': {
            'type': 'string'
        }
    }
}

mapreduce_noop = {
    'schema': {
        'name': {
            'type': 'string'
        },
        'description': {
            'type': 'string'
        }
    }
}

role = {
    'schema': {
        'users': {
            'type': 'list',
            'schema': {
                'type': 'objectid',
                'data_relation': {
                    'resource': 'user',
                    'field': '_id',
                    'embeddable': True
                }
            }
        },
        'name': {
            'type': 'string'
        },
        'description': {
            'type': 'string'
        }
    }
}
}

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

AzureNodeExtra = {
  'schema': {
    'ssh_port': {
      'type': 'string'
    },
    'instance_size': {
      'type': 'string'
    },
    'remote_desktop_port': {
      'type': 'string'
    },
    'ex_cloud_service_name': {
      'type': 'string'
    },
    'power_state': {
      'type': 'string'
    },
    'instance_endpoints': {
      'type': 'string'
    }
  }
}

var = {
  'schema': {
    'type': {
      'type': 'string'
    },
    'name': {
      'type': 'string'
    },
    'value': {
      'type': 'string'
    }
  }
}

profile = {
  'schema': {
    'username': {
      'type': 'string'
    },
    'context': {
      'type': 'string'
    },
    'description': {
      'type': 'string'
    },
    'firstname': {
      'type': 'string'
    },
    'lastname': {

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

        'type': 'string'
    },
    'publickey': {
        'type': 'string'
    },
    'email': {
        'type': 'string'
    },
    'uuid': {
        'type': 'string'
    }
}
}

virtual_machine = {
  'schema': {
    'status': {
      'type': 'string'
    },
    'authorized_keys': {
      'type': 'list',
      'schema': {
        'type': 'objectid',
        'data_relation': {
          'resource': 'sshkey',
          'field': '_id',
          'embeddable': True
        }
      }
    },
    'name': {
      'type': 'string'
    },
    'nics': {
      'type': 'list',
      'schema': {
        'type': 'objectid',
        'data_relation': {
          'resource': 'nic',
          'field': '_id',
          'embeddable': True
        }
      }
    },
    'RAM': {
      'type': 'string'
    },
    'ncpu': {
      'type': 'integer'
    },
    'loginuser': {
      'type': 'string'
    },
  },

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

        'disk': {
            'type': 'string'
        },
        'OS': {
            'type': 'string'
        },
        'metadata': {
            'type': 'dict',
            'schema': {}
        }
    }
}

kubernetes = {
    'schema': {
        'items': {
            'type': 'list',
            'schema': {
                'type': 'dict',
                'schema': {
                    'status': {
                        'type': 'dict',
                        'schema': {
                            'capacity': {
                                'type': 'dict',
                                'schema': {
                                    'cpu': {
                                        'type': 'string'
                                    }
                                }
                            }
                        }
                    },
                    'addresses': {
                        'type': 'list',
                        'schema': {
                            'type': 'dict',
                            'schema': {
                                'type': {
                                    'type': 'string'
                                },
                                'address': {
                                    'type': 'string'
                                }
                            }
                        }
                    }
                }
            }
        },
        'kind': {
            'type': 'string'
        },
        'metadata': {
            'type': 'dict',
            'schema': {}
        }
    }
}

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

        'name': {
            'type': 'string'
        }
    },
    'kind': {
        'type': 'string'
    },
    'users': {
        'type': 'list',
        'schema': {
            'type': 'dict',
            'schema': {
                'name': {
                    'type': 'string'
                },
                'user': {
                    'type': 'dict',
                    'schema': {
                        'username': {
                            'type': 'string'
                        },
                        'password': {
                            'type': 'string'
                        }
                    }
                }
            }
        }
    }
}

nic = {
    'schema': {
        'name': {
            'type': 'string'
        },
        'ip': {
            'type': 'string'
        },
        'mask': {
            'type': 'string'
        },
        'bandwidth': {
            'type': 'string'
        },
        'mtu': {
            'type': 'integer'
        }
    },

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>



**Object A.1: Schema**

```

        'broadcast': {
            'type': 'string'
        },
        'mac': {
            'type': 'string'
        },
        'type': {
            'type': 'string'
        },
        'gateway': {
            'type': 'string'
        }
    }
}

openstack_flavor = {
    'schema': {
        'os_flv_disabled': {
            'type': 'string'
        },
        'uuid': {
            'type': 'string'
        },
        'os_flv_ext_data': {
            'type': 'string'
        },
        'ram': {
            'type': 'string'
        },
        'os_flavor_acces': {
            'type': 'string'
        },
        'vcpus': {
            'type': 'string'
        },
        'swap': {
            'type': 'string'
        },
        'rxtx_factor': {
            'type': 'string'
        },
        'disk': {
            'type': 'string'
        }
    }
}

azure_vm = {
    'schema': {
        'username': {
            'type': 'string'
        },
        'status': {

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

        'type': 'string'
    },
    'deployment_slot': {
        'type': 'string'
    },
    'group': {
        'type': 'string'
    },
    'private_ips': {
        'type': 'string'
    },
    'cloud_service': {
        'type': 'string'
    },
    'dns_name': {
        'type': 'string'
    },
    'image': {
        'type': 'string'
    },
    'floating_ip': {
        'type': 'string'
    },
    'image_name': {
        'type': 'string'
    },
    'instance_name': {
        'type': 'string'
    },
    'public_ips': {
        'type': 'string'
    },
    'media_link': {
        'type': 'string'
    },
    'key': {
        'type': 'string'
    },
    'flavor': {
        'type': 'string'
    },
    'resource_location': {
        'type': 'string'
    },
    'instance_size': {
        'type': 'string'
    },
    'disk_name': {
        'type': 'string'
    },
    'uuid': {
        'type': 'string'
    }
}
    
```

**Object A.1: Schema**

```

    }
  }

  ec2NodeExtra = {
    'schema': {
      'product_codes': {
        'type': 'string'
      },
      'tags': {
        'type': 'list',
        'schema': {
          'type': 'string'
        }
      },
      'network_interfaces': {
        'type': 'list',
        'schema': {
          'type': 'string'
        }
      },
      'groups': {
        'type': 'list',
        'schema': {
          'type': 'string'
        }
      },
      'block_device_mapping': {
        'type': 'string'
      }
    }
  }

  libcloud_image = {
    'schema': {
      'username': {
        'type': 'string'
      },
      'status': {
        'type': 'string'
      },
      'updated': {
        'type': 'string'
      },
      'description': {
        'type': 'string'
      },
      'owner_alias': {
        'type': 'string'
      },
      'kernel_id': {
        'type': 'string'
      },
      'hypervisor': {

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

        'type': 'string'
    },
    'ramdisk_id': {
        'type': 'string'
    },
    'state': {
        'type': 'string'
    },
    'created': {
        'type': 'string'
    },
    'image_id': {
        'type': 'string'
    },
    'image_location': {
        'type': 'string'
    },
    'platform': {
        'type': 'string'
    },
    'image_type': {
        'type': 'string'
    },
    'is_public': {
        'type': 'string'
    },
    'owner_id': {
        'type': 'string'
    },
    'architecture': {
        'type': 'string'
    },
    'virtualization_type': {
        'type': 'string'
    },
    'uuid': {
        'type': 'string'
    }
}
}

user = {
    'schema': {
        'profile': {
            'type': 'objectid',
            'data_relation': {
                'resource': 'profile',
                'field': '_id',
                'embeddable': True
            }
        }
    },
    'roles': {
        'type': 'list',

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

        'schema': {
            'type': 'string'
        }
    }
}

```

```

GCENodeExtra = {
  'schema': {
    'status': {
      'type': 'string'
    },
    'kind': {
      'type': 'string'
    },
    'machineType': {
      'type': 'string'
    },
    'description': {
      'type': 'string'
    },
    'zone': {
      'type': 'string'
    },
    'deprecated': {
      'type': 'string'
    },
    'image': {
      'type': 'string'
    },
    'disks': {
      'type': 'string'
    },
    'tags_fingerprint': {
      'type': 'string'
    },
    'name': {
      'type': 'string'
    },
    'boot_disk': {
      'type': 'string'
    },
    'selfLink': {
      'type': 'string'
    },
    'scheduling': {
      'type': 'string'
    },
    'canIpForward': {
      'type': 'string'
    },
    'serviceAccounts': {
      'type': 'string'
    }
  }
}

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

    },
    'metadata': {
      'type': 'string'
    },
    'creationTimestamp': {
      'type': 'string'
    },
    'id': {
      'type': 'string'
    },
    'statusMessage': {
      'type': 'string'
    },
    'networkInterfaces': {
      'type': 'string'
    }
  }
}

group = {
  'schema': {
    'users': {
      'type': 'list',
      'schema': {
        'type': 'objectid',
        'data_relation': {
          'resource': 'user',
          'field': '_id',
          'embeddable': True
        }
      }
    },
    'name': {
      'type': 'string'
    },
    'description': {
      'type': 'string'
    }
  }
}

secgroup = {
  'schema': {
    'ingress': {
      'type': 'string'
    },
    'egress': {
      'type': 'string'
    },
    'ports': {
      'type': 'integer'
    },
    'protocols': {

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

        'type': 'string'
    }
}

node_new = {
    'schema': {
        'authorized_keys': {
            'type': 'list',
            'schema': {
                'type': 'string'
            }
        },
        'name': {
            'type': 'string'
        },
        'external_ip': {
            'type': 'string'
        },
        'memory': {
            'type': 'integer'
        },
        'create_external_ip': {
            'type': 'boolean'
        },
        'internal_ip': {
            'type': 'string'
        },
        'loginuser': {
            'type': 'string'
        },
        'owner': {
            'type': 'string'
        },
        'cores': {
            'type': 'integer'
        },
        'disk': {
            'type': 'integer'
        },
        'ssh_keys': {
            'type': 'list',
            'schema': {
                'type': 'dict',
                'schema': {
                    'from': {
                        'type': 'string'
                    },
                    'decrypt': {
                        'type': 'string'
                    },
                    'ssh_keygen': {
                        'type': 'boolean'
                    }
                }
            }
        }
    }
}

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

        },
        'to': {
            'type': 'string'
        }
    }
},
'security_groups': {
    'type': 'list',
    'schema': {
        'type': 'dict',
        'schema': {
            'ingress': {
                'type': 'string'
            },
            'egress': {
                'type': 'string'
            },
            'ports': {
                'type': 'list',
                'schema': {
                    'type': 'integer'
                }
            },
            'protocols': {
                'type': 'list',
                'schema': {
                    'type': 'string'
                }
            }
        }
    }
},
'users': {
    'type': 'dict',
    'schema': {
        'name': {
            'type': 'string'
        },
        'groups': {
            'type': 'list',
            'schema': {
                'type': 'string'
            }
        }
    }
}
}

batchjob = {
    'schema': {
        'output_file': {

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>



**Object A.1: Schema**

```

        'type': 'string'
    },
    'group': {
        'type': 'string'
    },
    'job_id': {
        'type': 'string'
    },
    'script': {
        'type': 'string'
    },
    'cmd': {
        'type': 'string'
    },
    'queue': {
        'type': 'string'
    },
    'cluster': {
        'type': 'string'
    },
    'time': {
        'type': 'string'
    },
    'path': {
        'type': 'string'
    },
    'nodes': {
        'type': 'string'
    },
    'dir': {
        'type': 'string'
    }
    }
}

account = {
    'schema': {
        'status': {
            'type': 'string'
        },
        'startDate': {
            'type': 'string'
        },
        'endDate': {
            'type': 'string'
        },
        'description': {
            'type': 'string'
        },
        'uuid': {
            'type': 'string'
        },
        'user': {

```

**Object A.1: Schema**

```

        'type': 'list',
        'schema': {
            'type': 'string'
        }
    },
    'group': {
        'type': 'list',
        'schema': {
            'type': 'string'
        }
    },
    'balance': {
        'type': 'float'
    },
    'name': {
        'type': 'string'
    }
}

libcloud_vm = {
    'schema': {
        'username': {
            'type': 'string'
        },
        'status': {
            'type': 'string'
        },
        'root_device_type': {
            'type': 'string'
        },
        'private_ips': {
            'type': 'string'
        },
        'instance_type': {
            'type': 'string'
        },
        'image': {
            'type': 'string'
        },
        'private_dns': {
            'type': 'string'
        },
        'image_name': {
            'type': 'string'
        },
        'instance_id': {
            'type': 'string'
        },
        'image_id': {
            'type': 'string'
        },
        'public_ips': {

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

        'type': 'string'
    },
    'state': {
        'type': 'string'
    },
    'root_device_name': {
        'type': 'string'
    },
    'key': {
        'type': 'string'
    },
    'group': {
        'type': 'string'
    },
    'flavor': {
        'type': 'string'
    },
    'availability': {
        'type': 'string'
    },
    'uuid': {
        'type': 'string'
    }
}
}

compute_node = {
    'schema': {
        'status': {
            'type': 'string'
        },
        'authorized_keys': {
            'type': 'list',
            'schema': {
                'type': 'objectid',
                'data_relation': {
                    'resource': 'sshkey',
                    'field': '_id',
                    'embeddable': True
                }
            }
        },
        'kind': {
            'type': 'string'
        },
        'uuid': {
            'type': 'string'
        },
        'secgroups': {
            'type': 'list',
            'schema': {
                'type': 'objectid',
                'data_relation': {

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

        'resource': 'secgroup',
        'field': '_id',
        'embeddable': True
    }
}
},
'nics': {
    'type': 'list',
    'schema': {
        'type': 'objectid',
        'data_relation': {
            'resource': 'nic',
            'field': '_id',
            'embeddable': True
        }
    }
},
'image': {
    'type': 'string'
},
'label': {
    'type': 'string'
},
'loginuser': {
    'type': 'string'
},
'flavor': {
    'type': 'list',
    'schema': {
        'type': 'objectid',
        'data_relation': {
            'resource': 'flavor',
            'field': '_id',
            'embeddable': True
        }
    }
},
'metadata': {
    'type': 'dict',
    'schema': {
        'owner': {
            'type': 'string'
        },
        'experiment': {
            'type': 'string'
        }
    }
},
'name': {
    'type': 'string'
}
}
}

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

database = {
  'schema': {
    'endpoint': {
      'type': 'string'
    },
    'protocol': {
      'type': 'string'
    },
    'name': {
      'type': 'string'
    }
  }
}

default = {
  'schema': {
    'context': {
      'type': 'string'
    },
    'name': {
      'type': 'string'
    },
    'value': {
      'type': 'string'
    }
  }
}

openstack_image = {
  'schema': {
    'status': {
      'type': 'string'
    },
    'username': {
      'type': 'string'
    },
    'updated': {
      'type': 'string'
    },
    'uuid': {
      'type': 'string'
    },
    'created': {
      'type': 'string'
    },
    'minDisk': {
      'type': 'string'
    },
    'progress': {
      'type': 'string'
    },
    'minRam': {

```

**Object A.1: Schema**

```

        'type': 'string'
    },
    'os_image_size': {
        'type': 'string'
    },
    'metadata': {
        'type': 'dict',
        'schema': {
            'instance_uuid': {
                'type': 'string'
            },
            'image_location': {
                'type': 'string'
            },
            'image_state': {
                'type': 'string'
            },
            'instance_type_memory_mb': {
                'type': 'string'
            },
            'user_id': {
                'type': 'string'
            },
            'description': {
                'type': 'string'
            },
            'kernel_id': {
                'type': 'string'
            },
            'instance_type_name': {
                'type': 'string'
            },
            'ramdisk_id': {
                'type': 'string'
            },
            'instance_type_id': {
                'type': 'string'
            },
            'instance_type_ephemeral_gb': {
                'type': 'string'
            },
            'instance_type_rxtx_factor': {
                'type': 'string'
            },
            'image_type': {
                'type': 'string'
            },
            'network_allocated': {
                'type': 'string'
            },
            'instance_type_flavorid': {
                'type': 'string'
            },
        },
    },

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```
        'instance_type_vcpus': {
            'type': 'string'
        },
        'instance_type_root_gb': {
            'type': 'string'
        },
        'base_image_ref': {
            'type': 'string'
        },
        'instance_type_swap': {
            'type': 'string'
        },
        'owner_id': {
            'type': 'string'
        }
    }
}
}
}

OpenStackNodeExtra = {
  'schema': {
    'vm_state': {
      'type': 'string'
    },
    'addresses': {
      'type': 'list',
      'schema': {
        'type': 'string'
      }
    },
    'availability_zone': {
      'type': 'string'
    },
    'service_name': {
      'type': 'string'
    },
    'userId': {
      'type': 'string'
    },
    'imageId': {
      'type': 'string'
    },
    'volumes_attached': {
      'type': 'string'
    },
    'task_state': {
      'type': 'string'
    },
    'disk_config': {
      'type': 'string'
    },
    'power_state': {
```

**Object A.1: Schema**

```

        'type': 'string'
    },
    'progress': {
        'type': 'string'
    },
    'metadata': {
        'type': 'list',
        'schema': {
            'type': 'string'
        }
    },
    'updated': {
        'type': 'string'
    },
    'hostId': {
        'type': 'string'
    },
    'key_name': {
        'type': 'string'
    },
    'flavorId': {
        'type': 'string'
    },
    'password': {
        'type': 'string'
    },
    'access_ip': {
        'type': 'string'
    },
    'access_ipv6': {
        'type': 'string'
    },
    'created': {
        'type': 'string'
    },
    'fault': {
        'type': 'string'
    },
    'uri': {
        'type': 'string'
    },
    'tenantId': {
        'type': 'string'
    },
    'config_drive': {
        'type': 'string'
    }
}

mapreduce = {
    'schema': {
        'function': {

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>



**Object A.1: Schema**

```

        'type': 'dict',
        'schema': {
            'source': {
                'type': 'string'
            },
            'args': {
                'type': 'dict',
                'schema': {}
            }
        }
    },
    'fault_tolerant': {
        'type': 'boolean'
    },
    'data': {
        'type': 'dict',
        'schema': {
            'dest': {
                'type': 'string'
            },
            'source': {
                'type': 'string'
            }
        }
    },
    'backend': {
        'type': 'dict',
        'schema': {
            'type': {
                'type': 'string'
            }
        }
    }
}

filter = {
    'schema': {
        'function': {
            'type': 'string'
        },
        'name': {
            'type': 'string'
        }
    }
}

alias = {
    'schema': {
        'origin': {
            'type': 'string'
        },
        'name': {

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

        'type': 'string'
    }
}

replica = {
  'schema': {
    'endpoint': {
      'type': 'string'
    },
    'name': {
      'type': 'string'
    },
    'checksum': {
      'type': 'dict',
      'schema': {
        'md5': {
          'type': 'string'
        }
      }
    },
    'replica': {
      'type': 'string'
    },
    'accessed': {
      'type': 'string'
    },
    'size': {
      'type': 'list',
      'schema': {
        'type': 'string'
      }
    }
  }
}

openstack_vm = {
  'schema': {
    'vm_state': {
      'type': 'string'
    },
    'availability_zone': {
      'type': 'string'
    },
    'terminated_at': {
      'type': 'string'
    },
    'image': {
      'type': 'string'
    },
    'diskConfig': {
      'type': 'string'
    },
  },
}

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

    'flavor': {
      'type': 'string'
    },
    'security_groups': {
      'type': 'string'
    },
    'volumes_attached': {
      'type': 'string'
    },
    'user_id': {
      'type': 'string'
    },
    'uuid': {
      'type': 'string'
    },
    'accessIPv4': {
      'type': 'string'
    },
    'accessIPv6': {
      'type': 'string'
    },
    'power_state': {
      'type': 'string'
    },
    'progress': {
      'type': 'string'
    },
    'image_id': {
      'type': 'string'
    },
    'launched_at': {
      'type': 'string'
    },
    'config_drive': {
      'type': 'string'
    },
    'username': {
      'type': 'string'
    },
    'updated': {
      'type': 'string'
    },
    'hostId': {
      'type': 'string'
    },
    'floating_ip': {
      'type': 'string'
    },
    'static_ip': {
      'type': 'string'
    },
    'key': {
      'type': 'string'
    }

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

    },
    'flavor__id': {
        'type': 'string'
    },
    'group': {
        'type': 'string'
    },
    'task_state': {
        'type': 'string'
    },
    'created': {
        'type': 'string'
    },
    'tenant_id': {
        'type': 'string'
    },
    'status': {
        'type': 'string'
    }
}
}

organization = {
    'schema': {
        'users': {
            'type': 'list',
            'schema': {
                'type': 'objectid',
                'data_relation': {
                    'resource': 'user',
                    'field': '_id',
                    'embeddable': True
                }
            }
        }
    }
}

hadoop = {
    'schema': {
        'deployers': {
            'type': 'dict',
            'schema': {
                'ansible': {
                    'type': 'string'
                }
            }
        },
        'requires': {
            'type': 'dict',
            'schema': {
                'java': {
                    'type': 'dict',

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

        'schema': {
            'implementation': {
                'type': 'string'
            },
            'version': {
                'type': 'string'
            },
            'zookeeper': {
                'type': 'string'
            },
            'supervisord': {
                'type': 'string'
            }
        }
    },
    'parameters': {
        'type': 'dict',
        'schema': {
            'num_resourcemangers': {
                'type': 'integer'
            },
            'num_namenodes': {
                'type': 'integer'
            },
            'use_yarn': {
                'type': 'boolean'
            },
            'num_datanodes': {
                'type': 'integer'
            },
            'use_hdfs': {
                'type': 'boolean'
            },
            'num_historyservers': {
                'type': 'integer'
            },
            'num_journalnodes': {
                'type': 'integer'
            }
        }
    }
}

accounting_resource = {
    'schema': {
        'account': {
            'type': 'string'
        },
        'group': {
            'type': 'string'
        }
    }
}

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

**Object A.1: Schema**

```

    },
    'description': {
        'type': 'string'
    },
    'parameters': {
        'type': 'dict',
        'schema': {
            'parameter1': {
                'type': 'float'
            },
            'parameter2': {
                'type': 'float'
            }
        }
    },
    'uuid': {
        'type': 'string'
    },
    'charge': {
        'type': 'string'
    },
    'unites': {
        'type': 'dict',
        'schema': {
            'parameter1': {
                'type': 'string'
            },
            'parameter2': {
                'type': 'string'
            }
        }
    },
    'user': {
        'type': 'string'
    },
    'name': {
        'type': 'string'
    }
}

eve_settings = {
    'MONGO_HOST': 'localhost',
    'MONGO_DBNAME': 'testing',
    'RESOURCE_METHODS': ['GET', 'POST', 'DELETE'],
    'BANDWIDTH_SAVER': False,
    'DOMAIN': {
        'container': container,
        'stream': stream,
        'azure_image': azure_image,
        'deployment': deployment,
    }
}

```

**Object A.1: Schema**

```

    'azure-size': azure_size,
    'cluster': cluster,
    'computer': computer,
    'mesos-docker': mesos_docker,
    'file': file,
    'reservation': reservation,
    'microservice': microservice,
    'flavor': flavor,
    'virtual_directory': virtual_directory,
    'mapreduce_function': mapreduce_function,
    'virtual_cluster': virtual_cluster,
    'libcloud_flavor': libcloud_flavor,
    'LibCloudNode': LibCloudNode,
    'sshkey': sshkey,
    'timestamp': timestamp,
    'mapreduce_noop': mapreduce_noop,
    'role': role,
    'AzureNodeExtra': AzureNodeExtra,
    'var': var,
    'profile': profile,
    'virtual_machine': virtual_machine,
    'kubernetes': kubernetes,
    'nic': nic,
    'openstack_flavor': openstack_flavor,
    'azure-vm': azure_vm,
    'ec2NodeExtra': ec2NodeExtra,
    'libcloud_image': libcloud_image,
    'user': user,
    'GCENodeExtra': GCENodeExtra,
    'group': group,
    'secgroup': secgroup,
    'node_new': node_new,
    'batchjob': batchjob,
    'account': account,
    'libcloud_vm': libcloud_vm,
    'compute_node': compute_node,
    'database': database,
    'default': default,
    'openstack_image': openstack_image,
    'OpenStackNodeExtra': OpenStackNodeExtra,
    'mapreduce': mapreduce,
    'filter': filter,
    'alias': alias,
    'replica': replica,
    'openstack_vm': openstack_vm,
    'organization': organization,
    'hadoop': hadoop,
    'accounting_resource': accounting_resource,
},
}

```

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

# Appendix B: Cloudmesh REST

---

Cloudmesh Rest is a reference implementation for the NBDRA. It allows for automatic definition of a REST service based on the objects specified by the NBDRA. In collaboration with other cloudmesh components, it allows easy interaction with hybrid clouds and the creation of user-managed Big Data services.

## B.1 Prerequisites

The prerequisites for cloudmesh Rest are Python 2.7.13 or 3.6.1. It can easily be installed on a variety of systems (At this time, only ubuntu greater 16.04 and OSX Sierra have been tested.). However, it would naturally be possible to also port it to Windows. At the time of publication, the installation instructions in this document are not complete. The reader is referred to the cloudmesh manuals, which are under development. The goal will be to make the installation (after the system is set up for developing Python) as simple as the following:

```
pip install cloudmesh.rest
```

## B.2 REST Service

With the cloudmesh REST framework, it is easy to create REST services while defining the resources via example JSON objects. This is achieved while leveraging the Python eve [13] and a modified version of Python evengine [14].

A valid JSON resource specification looks like this:

```
{
  "profile": {
    "description": "The Profile of a user",
    "email": "laszewski@gmail.com",
    "firstname": "Gregor",
    "lastname": "von Laszewski",
    "username": "gregor"
  }
}
```

In this example, an object called profile is defined, which contains a number of attributes and values. The type of the values is automatically determined. All JSON specifications are contained in a directory and can easily be converted into a valid schema for the eve REST service by executing the following commands:

```
cms schema cat . all.json
cms schema convert all.json
```

This will create the configuration `\verb|all.settings.py|` that can be used to start an eve service.

Once the schema has been defined, cloudmesh specifies defaults for managing a sample database that is coupled with the REST service. MongoDB was used, which could be placed on a shared mongo service.



## ***B.3 Limitations***

The current implementation is a demonstration that showcases the generation of a fully functioning REST service based on the specifications provided in this document. However, it is expected that scalability, distribution of services, and other advanced options need to be addressed based on application requirements.

# Appendix C: Acronyms and Terms

---

The following acronyms and terms are used in this volume.

ACID	Atomicity, Consistency, Isolation, Durability
API	Application Programming Interface
ASCII	American Standard Code for Information Interchange
BASE	Basically Available, Soft state, Eventual consistency
Container	See <a href="http://csrc.nist.gov/publications/drafts/800-180/sp800-180_draft.pdf">http://csrc.nist.gov/publications/drafts/800-180/sp800-180_draft.pdf</a>
Cloud Computing	The practice of using a network of remote servers hosted on the Internet to store, manage, and process data, rather than a local server or a personal computer. See <a href="http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf">http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf</a> .
DevOps	A clipped compound of software DEvelopment and information technology OPerationS
Deployment	The action of installing software on resources
HTTP	HyperText Transfer Protocol HTTPS HTTP Secure
Hybrid Cloud	See <a href="http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf">http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf</a> .
IaaS	Infrastructure as a Service SaaS Software as a Service
ITL	Information Technology Laboratory
Microservice Architecture	Is an approach to build applications based on many smaller modular services. Each module supports a specific goal and uses a simple, well-defined interface to communicate with other sets of services.
NBD-PWG	NIST Big Data Public Working Group
NBDRA	NIST Big Data Reference Architecture
NBDRAI	NIST Big Data Reference Architecture Interface
NIST	National Institute of Standards and Technology
OS	Operating System
REST	REpresentational State Transfer
Replica	A duplicate of a file on another resource to avoid costly transfer costs in case of frequent access.
Serverless Computing	Serverless computing specifies the paradigm of function as a service (FaaS). It is a cloud computing code execution model in which a cloud provider manages the function deployment and utilization while clients can utilize them. The charge model is based on execution of the function rather than the cost to manage and host the VM or container.
Software Stack	A set of programs and services that are installed on a resource to support applications.
Virtual Filesystem	An abstraction layer on top of a distributed physical file system to allow easy access to the files by the user or application.

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

Virtual Machine	A VM is a software computer that, like a physical computer, runs an operating system and applications. The VM is composed of a set of specification and configuration files and is backed by the physical resources of a host.
Virtual Cluster	A virtual cluster is a software cluster that integrate either VMs, containers, or physical resources into an agglomeration of compute resources. A virtual cluster allows users to authenticate and authorize to the virtual compute nodes to utilize them for calculations. Optional high-level services that can be deployed on a virtual cluster may simplify interaction with the virtual cluster or provide higher-level services.
Workflow	The sequence of processes or tasks
WWW	World Wide Web

This publication is available free of charge from: <https://doi.org/10.6028/NIST.SP.1500-9>

# Appendix D: Bibliography

---

- [1] W. Chang and NIST Big Data Public Working Group, “NIST Big Data Interoperability Framework: Volume 1, Definitions (SP1500-1),” 2015.
- [2] W. Chang and NIST Big Data Public Working Group, “NIST Big Data Interoperability Framework: Volume 2, Big Data Taxonomies (SP1500-2),” 2015.
- [3] W. Chang and NIST Big Data Public Working Group, “NIST Big Data Interoperability Framework: Volume 3, Use Cases and General Requirements (SP1500-3),” 2015.
- [4] W. Chang and NIST Big Data Public Working Group, “NIST Big Data Interoperability Framework: Volume 4, Security and Privacy (SP1500-4),” 2015.
- [5] W. Chang and NIST Big Data Public Working Group, “NIST Big Data Interoperability Framework: Volume 5, Architectures White Paper Survey (SP1500-5),” 2015.
- [6] W. Chang and NIST Big Data Public Working Group, “NIST Big Data Interoperability Framework: Volume 6, Reference Architecture (SP1500-6),” 2015.
- [7] W. Chang and NIST Big Data Public Working Group, “NIST Big Data Interoperability Framework: Volume 7, Standards Roadmap (SP1500-7),” 2015.
- [8] W. Chang and NIST Big Data Public Working Group, “NIST Big Data Interoperability Framework: Volume 9, Adoption and Modernization (SP1500-10),” 2017.
- [9] G. von Laszewski, F. Wang, B. Abdul-Wahid, H. Lee, G. C. Fox, and W. Chang, “DRAFT Cloudmesh in support of the NIST Big Data Architecture Framework,” Bloomington, IN, 2017.
- [10] G. von Laszewski, “Cloudmesh Client,” *GitHub*. [Online]. Available: <https://github.com/cloudmesh/client>. [Accessed: 15-Nov-2017].
- [11] N. Iarocci, “Welcome to Cerberus,” *Cerberus website*. [Online]. Available: <http://docs.python-cerberus.org/en/stable/>. [Accessed: 15-Nov-2017].
- [12] Internet2 Middleware Architecture Committee for Education Directory Working Group (MACE-Dir), “eduPerson Object Class Specification (201602),” 2016. [Online]. Available: <http://software.internet2.edu/eduperson/internet2-mace-dir-eduperson-201602.html>. [Accessed: 15-Nov-2017].
- [13] N. Iarocci, “Eve. The Simple Way to REST,” *Eve website*. [Online]. Available: <http://python-eve.org/>. [Accessed: 15-Nov-2017].
- [14] G. von Laszewski, “Cloudmesh Enhanced Eve Genie,” *GitHub*. [Online]. Available: <https://github.com/cloudmesh/cloudmesh.evegenie>. [Accessed: 15-Nov-2017].

---

<sup>1</sup> The White House Office of Science and Technology Policy, “Big Data is a Big Deal,” *OSTP Blog*, accessed February 21, 2014, <http://www.whitehouse.gov/blog/2012/03/29/big-data-big-deal>.